Chapter 3: MHD stability, operational limits and disruptions

2007 Nucl. Fusion 47 S128

(http://iopscience.iop.org/0029-5515/47/6/S03)

Nucl. Fusion **47** (2007) S128–S202 doi:10.1088/0029-5515/47/6/S03

Chapter 3: MHD stability, operational limits and disruptions

T.C. Hender[1,a], J.C Wesley[2], J. Bialek[3], A. Bondeson[4], A.H. Boozer[3], R.J. Buttery[1], A. Garofalo[3], T.P Goodman[5], R.S. Granetz[6], Y. Gribov[7], O. Gruber[8], M. Gryaznevich[1], G. Giruzzi[9], S. Günter[8], N. Hayashi[10], P. Helander[11], C.C. Hegna[12], D.F. Howell[1], D.A. Humphreys[2], G.T.A. Huysmans[9], A.W. Hyatt[2], A. Isayama[10], S.C. Jardin[13], Y. Kawano[10], A. Kellman[2], C. Kessel[13], H.R. Koslowski[14], R.J. La Haye[2], E. Lazzaro[15], Y.Q. Liu[4], V. Lukash[16], J. Manickam[13], S. Medvedev[17], V. Mertens[8], S.V. Mirnov[18], Y. Nakamura[10], G. Navratil[3], M. Okabayashi[13], T. Ozeki[10], R. Paccagnella[19], G. Pautasso[8], F. Porcelli[20], V.D. Pustovitov[16], V. Riccardo[1], M. Sato[10], O. Sauter[5], M.J. Schaffer[2], M. Shimada[7], P. Sonato[19], E.J. Strait[2], M. Sugihara[7], M. Takechi[10], A.D. Turnbull[2], E. Westerhof[21], D.G. Whyte[12], R. Yoshino[10], H. Zohm[8] and the ITPA MHD, Disruption and Magnetic Control Topical Group

[1] EURATOM/UKAEA Fusion Association, Culham Science Centre, Abingdon, Oxon, OX14 3DB, UK
[2] General Atomics, PO Box 85608, San Diego, California 92186-5608, USA
[3] Columbia University, New York, New York 10027, USA
[4] EURATOM/VR Fusion Association, Chalmers University of Technology, Göteborg, Sweden
[5] Ecole Polytechnique Fédérale de Lausanne (EPFL), Centre de Recherches en Physique des Plasmas, Association Euratom-Confédération Suisse, CH-1015 Lausanne, Switzerland
[6] Massachusetts Institute of Technology, Plasma Science and Fusion Center, Cambridge, MA 02139, USA
[7] ITER Organization, Cadarache Centre, 13108 St Paul lez Durance, France
[8] Max-Planck Institut für Plasmaphysik, EURATOM Association, D-85748 Garching, Germany
[9] Association Euratom-CEA, CEA Cadarache, F-13108 St Paul-lez-Durance, France
[10] Japan Atomic Energy Agency, Naka, Ibaraki-ken, 311-0193, Japan
[11] Max-Planck Institut für Plasmaphysik, EURATOM Association, D-17491 Greifswald, Germany
[12] Department of Engineering Physics, University of Wisconsin-Madison, Madison, Wisconsin 53706, USA
[13] Princeton Plasma Physics Laboratory, PO Box 451, Princeton, NJ 08543, USA
[14] Forschungszentrum Jülich GmbH, Association EURATOM-FZ Jülich, Institut für Plasmaphysik, Trilateral Euregio Cluster, D-52425 Jülich, Germany
[15] IFP CNR, Assoc. Euratom-ENEA-CNR, Via R. Cozzi 53, Milan, Italy
[16] Nuclear Fusion Institute, Russian Research Centre 'Kurchatov Institute', Moscow, Russian Federation
[17] Keldysh Institute of Applied Mathematics, Moscow, Russian Federation
[18] State Research Centre TRINITI, Troitsk, Moscow Region, 142190 Russian Federation
[19] Consorzio RFX, Assoc. Euratom-ENEA, I-35127 Padova, Italy
[20] Burning Plasma Research Group, Politecnico di Torino, Torino 10129, Italy
[21] FOM-Institute for Plasma Physics Rijnhuizen, Association EURATOM-FOM, Trilateral Euregio Cluster, PO Box 1207, 3430 BE Nieuwegein, Netherlands

E-mail: tim.hender@ukaea.org.uk

Received 20 June 2006, accepted for publication 8 January 2007
Published 1 June 2007
Online at stacks.iop.org/NF/47/S128

Abstract
Progress in the area of MHD stability and disruptions, since the publication of the 1999 ITER Physics Basis document (1999 *Nucl. Fusion* **39** 2137–2664), is reviewed. Recent theoretical and experimental research has made important advances in both understanding and control of MHD stability in tokamak plasmas. Sawteeth are anticipated in the ITER baseline ELMy H-mode scenario, but the tools exist to avoid or control them through localized current drive or fast ion generation. Active control of other MHD instabilities will most likely be also required in ITER. Extrapolation from existing experiments indicates that stabilization of neoclassical tearing modes by highly localized feedback-controlled current drive should be possible in ITER. Resistive wall modes are a key issue for

[a] Author to whom any correspondence should be addressed.

advanced scenarios, but again, existing experiments indicate that these modes can be stabilized by a combination of plasma rotation and direct feedback control with non-axisymmetric coils. Reduction of error fields is a requirement for avoiding non-rotating magnetic island formation and for maintaining plasma rotation to help stabilize resistive wall modes. Recent experiments have shown the feasibility of reducing error fields to an acceptable level by means of non-axisymmetric coils, possibly controlled by feedback. The MHD stability limits associated with advanced scenarios are becoming well understood theoretically, and can be extended by tailoring of the pressure and current density profiles as well as by other techniques mentioned here. There have been significant advances also in the control of disruptions, most notably by injection of massive quantities of gas, leading to reduced halo current fractions and a larger fraction of the total thermal and magnetic energy dissipated by radiation. These advances in disruption control are supported by the development of means to predict impending disruption, most notably using neural networks. In addition to these advances in means to control or ameliorate the consequences of MHD instabilities, there has been significant progress in improving physics understanding and modelling. This progress has been in areas including the mechanisms governing NTM growth and seeding, in understanding the damping controlling RWM stability and in modelling RWM feedback schemes. For disruptions there has been continued progress on the instability mechanisms that underlie various classes of disruption, on the detailed modelling of halo currents and forces and in refining predictions of quench rates and disruption power loads. Overall the studies reviewed in this chapter demonstrate that MHD instabilities can be controlled, avoided or ameliorated to the extent that they should not compromise ITER operation, though they will necessarily impose a range of constraints.

PACS numbers: 52.55.Tn, 52.55.Fa

(Some figures in this article are in colour only in the electronic version)

Contents

1. Introduction
2. MHD stability
   - 2.1. Sawtooth oscillations
     - 2.1.1. Physics of sawtooth oscillations
     - 2.1.2. Sawtooth control
     - 2.1.3. Central MHD activity expected in ITER
   - 2.2. Neoclassical tearing modes
     - 2.2.1. Physics of neoclassical tearing modes
     - 2.2.2. Active control of NTMs
     - 2.2.3. Mitigation of NTMs
     - 2.2.4. NTMs expected in ITER and their control
   - 2.3. Resistive wall modes
     - 2.3.1. Physics of resistive wall modes
     - 2.3.2. Control of RWMs
     - 2.3.3. RWMs expected in ITER and their control
   - 2.4. Error fields
     - 2.4.1. Effect of error fields on plasmas
     - 2.4.2. Error fields expected in ITER and their correction
   - 2.5. MHD stability in advanced scenarios
     - 2.5.1. MHD stability in plasmas with weak magnetic shear
     - 2.5.2. MHD stability in plasmas with strong negative magnetic shear
     - 2.5.3. Expectations for ITER
   - 2.6. Summary
3. Disruptions
   - 3.1. Disruption characteristics, causes and frequency
   - 3.2. Thermal quench energy loss and deposition
   - 3.3. Current quench dynamics
     - 3.3.1. Global toroidal current decay
     - 3.3.2. Vertical instability, halo currents and mechanical forces
   - 3.4. Runaway electrons generated by disruptions
     - 3.4.1. Observations in present tokamaks
     - 3.4.2. Interaction of runaway electrons with plasma facing components
     - 3.4.3. Confinement, termination and mitigation of runaway electrons
     - 3.4.4. Summary and implications for ITER

3.5. Integrated modelling and simulation
3.5.1. ITER disruption simulations
3.5.2. DINA disruption modelling basis
3.5.3. Predictions for major disruptions and VDEs in ITER
3.5.4. Calculation of non-axisymmetric halo current effects in ITER by M3D code
3.5.5. Summary and future needs
3.6. Disruption avoidance, prediction and mitigation
3.6.1. Disruption prediction
Discussion and future work
3.6.2. Disruption avoidance
Passive avoidance
Disruption avoidance strategies
3.6.3. Disruption mitigation
Methods of disruption mitigation
Pellet injection
Massive gas injection
Liquid jets
Other concepts
Application to ITER
3.7. Summary and R&D needs

1. Introduction

Magnetohydrodynamic (MHD) stability and the microscopic and macroscopic effects of various classes of MHD instability underlie essentially all aspects of achievable plasma performance in tokamaks and determine the principal operational limits for tokamaks—maximum plasma current and plasma pressure (beta) and pressure gradient—and plays an important role in the closely related operational limit on maximum achievable or usable plasma density. The beta and density limits are two of the three key physics basis considerations—energy confinement is the third—that govern the basic design and plasma performance of reactor tokamaks. The combination of these three considerations determines the achievable fusion power and neutron wall loading and fusion power gain ($Q = P_{\rm fus}/P_{\rm aux}$; with $P_{\rm fus}$ being the fusion power and $P_{\rm aux}$ the auxiliary power heating of the plasma).

In contrast, disruptions, which arise from the immediate or eventual consequences of MHD instability, set a second type of constraint on the operational feasibility of reactor tokamaks—they can determine the operational lifetime of certain components, especially those associated with plasma power and particle exhaust. Since occurrence of at least a small number of disruptions is inevitable in any tokamak that seeks to explore burning plasma science in the reactor regime, the at-risk components and systems—especially the plasma-facing-component substrates and surfaces—*must* be designed to accommodate the immediate effects of a finite number of disruptions without premature failure or loss of function. Furthermore, since the consequences of disruption will likely shorten the usable operation life of plasma-facing-components, it is highly desirable to (1) minimize both the number and severity of disruptions that do occur and (2) ameliorate or 'soften' the consequences of disruptions that cannot otherwise be avoided. To this end, reactor tokamaks and ITER require (1) means to reliably predict the pending occurrence of disruption and (2) provision to couple these disruption predictions to one or more means for 'fast-shutdown' intervention which (ideally) will make it possible to either avoid the occurrence of disruption or to mitigate or soften the consequences of those disruptions which cannot be avoided.

A comprehensive discussion of the various aspects of MHD stability that contribute to the limits on plasma current and beta and pressure gradient anticipated in ITER, together with a discussion of the specialized nomenclature that applies to consideration of MHD stability in tokamak systems, can be found in chapter 3 of the ITER Physics Basis (IPB) [1]. The discussion therein proceeds from a fundamental understanding of ideal and resistive MHD stability theory to the more specific manifestations and consequences (operational limits) of various major categories of ideal and/or resistive MHD instabilities: e.g. sawtooth oscillations, neoclassical tearing modes, error-field-induced locked tearing modes, edge-localized modes (ELMs) and fast-growing internal modes localized near regions of high internal pressure gradient in advanced-performance plasmas with internal transport barriers (ITBs).

The MHD stability content of this chapter 3 of this issue 'Progress in the ITER Physics Basis' follows the same format and approach as the IPB chapter 3. Section 2.1 discusses the present understanding of the physics, control and application to ITER of the $m = 1, n = 1$ internal reconnection instability localized near the $q = 1$ surface that commonly manifests itself as sawtooth oscillations. The principal concern about sawtooth oscillations in ITER lies in their potential to provide 'seed islands' (due to the occurrence of a radial magnetic field resonant at given $q = m/n$ surfaces) that in turn can lead to onset of neoclassical tearing mode (NTM) instability. Section 2.2 addresses the physics, avoidance, control and effect on ITER of this resistive MHD instability. In ITER, uncontrolled growth of certain classes of NTM instability is predicted to lead to deterioration of energy confinement and possibly, at high-enough plasma pressure, to the onset of a disruption. The ITER 'reference case' ELMy H-mode

plasma operation scenario, with a positive magnetic shear profile, is particularly susceptible to various NTMs. The present understanding of the physics for the onset of NTMs in ITER suggests that NTMs in positive-shear scenarios will be inevitable if large enough seed instabilities (e.g. sawteeth) occur. Thus as Section 2.2.4 details, the present ITER design includes provision of electron cyclotron current drive (ECCD) for NTM control which is highly desirable for ensuring that ITER is able to reach its sustained burning plasma operation goal using the ELMy H-mode scenario.

Section 2.3 addresses the broad topic of resistive wall modes (RWMs), which are expected to be the principal beta-limiting instability in 'advanced-performance' steady-state-capable plasmas, based on 'reversed shear' or 'negative shear' operation. Resistive wall modes, which arise in any tokamak plasma when the plasma pressure exceeds the ideal $n = 1$ external kink 'no-wall' beta limit, are expected to set the beta limit (and disruption-onset beta) in ITER steady-state scenarios. As section 2.3 explains, RWMs can be controlled either by plasma rotation or by providing feedback control of the mode's non-axisymmetric field, or by a combination of the two. Section 2.3.3 details the critical-issue for advanced tokamak (AT) scenarios of how rotational and feedback stabilization will affect ITER's capability for AT steady-state operation.

The effects of small departures of the ITER magnetic field configuration from exact axisymmetry is addressed in section 2.4. A strategy for ITER that identifies construction and alignment accuracies for the toroidal and poloidal field coil systems, and the provision of a multi-mode 'error field correction' coil system is expected to allow successful mitigation of error-field-induced instabilities and disruptions. Such a practice is now routine in many present generation tokamaks.

The effect of localized internal MHD instabilities excited by excessive pressure gradients associated with ITBs in ITER plasmas is addressed in section 2.5. The principal conclusion here is that to avoid such instabilities becoming the limiting factor during ITER operation with ITB plasmas (as is envisioned for the reference case ITER steady-state scenario), undue localization of the barrier must be avoided. This requirement may, in turn, mean that the ITB pressure gradient of such a plasma must be controlled. This subject of ITB pressure gradient control is addressed in chapter 6 of this issue [2].

Section 2.6 provides a summary and synthesis of the prospects for achieving the MHD stability and the control of instabilities needed for successful operation in the various scenarios envisioned for ITER. 'Critical path' R&D needs in each of the MHD topical areas are also identified here.

The related subjects of edge-localized modes (ELMs) and density limit associated with the plasma boundary power balance are treated in chapter 4 of this issue [3], while that of energetic-particle-modes (toroidal Alfvén eigenmodes (TAE) and the like) is treated in chapter 6 of this issue [2].

The remainder of the chapter (section 3) covers disruption physics and characteristics, disruption modelling and simulation and disruption avoidance, prediction and mitigation. Again, the organization and nomenclature largely follows the corresponding section 3.4 of the 1999 ITER Physics Basis. Here, section 3.1 provides an introduction to disruption phenomenology and characteristics, and an overview of how these disruptions impact on various ITER systems and components. Sections 3.2 and 3.3 address details of the thermal quench and current quench phases of a disruption, respectively. Section 3.3 also includes discussion of the effect of vertical instability (the so-called vertical displacement event or VDE) and the resulting generation of circulating poloidal currents (halo currents) in the ITER blanket and vacuum vessel structures. Section 3.4 addresses the critical area of runaway electron generation in an ITER-class tokamak and the means/prospects for runaway avoidance/mitigation.

Section 3.5 addresses the status of modelling of disruptions and VDEs, and of halo currents and runaway conversion in ITER. Section 3.6 explores the closely related subjects of disruption prediction and the prospects for taking action to avoid a pending disruption or to implement a fast-shutdown action to avoid or at least partially mitigate the most serious consequences of disruption.

Section 3.7 concludes the disruption part of the chapter with a summary and synthesis of the expected ITER disruption characteristics and prospects for achieving the disruption prediction, avoidance and disruption effect mitigation needed for successful operation of ITER. 'Critical path' R&D needs in each of the disruption characterization and prediction/avoidance/mitigation topical areas are also identified here.

2. MHD stability

2.1. Sawtooth oscillations

2.1.1. Physics of sawtooth oscillations. Sawtooth oscillations are periodic, MHD initiated mixing events that occur in a tokamak plasma in the near axis region where the safety factor q is less than or equal to unity. Small sawteeth, with mixing radii of about 20–40% of the plasma minor radius and temperature drops of the order of a fraction of a keV, can be survived easily by the plasma discharge. Indeed, small sawteeth may even be beneficial in preventing the accumulation of impurities and helium ash in the plasma centre. On the other hand, large sawteeth with mixing radii of 50% or more of the plasma minor radius and temperature drops of one or more keV represent a serious threat to ITER operation, since these large sawtooth events may couple to NTMs [4] and to edge-localized modes [5], resulting in a serious loss of plasma energy and confinement degradation. An additional worry, in a fusion burning plasma, is that if the sawtooth period is shorter than the slowing down time of the fusion alpha particles then fusion alphas may be scattered, and perhaps lost, before they have time to transfer their energy to the thermal plasma. Therefore, the best sawtooth regime for ITER would be one where the sawtooth period is intermediate between the alpha particle slowing down time and the plasma heating (or energy confinement) time. Thus, in recent years a large effort has been dedicated to developing techniques for sawtooth control, reviewed in section 2.1.2, and in predicting the sawtooth behaviour in projected ITER discharges, as discussed in section 2.1.3. In this section, we review briefly the present theoretical understanding of the sawtooth phenomenon and indicate areas where further theoretical work is needed.

The underlying mechanism responsible for the sawtooth crash is MHD instability of an internal kink mode, with toroidal mode number $n = 1$ and dominant poloidal mode number $m = 1$. However, it is clear from experiments that ideal MHD theory alone is not accurate enough in predicting the threshold for the onset of the sawtooth crash. Non-ideal effects, such as kinetic effects associated with energetic particles (including the fusion alphas), finite electrical resistivity and viscosity, diamagnetic effects, finite ion Larmor radius, electron inertia, electron compressibility, play an important role in determining the actual stability threshold for the sawtooth crash. These effects have been incorporated in a phenomenological sawtooth trigger criterion, first proposed in [6] (see also section 2.1.3, where the model is described in some detail), and the reader is referred to this reference and to other references cited therein for a discussion of the non-ideal physics that determines the actual linear stability threshold of $m = n = 1$ modes in realistic tokamak discharges. This threshold depends on the actual collisionality regime (i.e. collisional, semi-collisional or collisionless), on the distribution of energetic ions in phase space and on the radial profiles of the plasma current density and pressure, including their local gradients near the $q = 1$ surface. Thus theoretical studies suggest that an efficient way to control sawteeth would be by affecting the value of magnetic shear near the $q = 1$ surface, for instance by means of localized heating and/or current drive.

The sawtooth trigger criterion has been incorporated in transport codes and has proved to be relatively successful in predicting quantitatively the sawtooth period and amplitude in existing tokamak experiments, as discussed in section 2.1.2. However, it is important to remark that this trigger criterion has a number of limitations. First of all, being predicated on linear theory, it may not be adequate to describe situations where long-lived precursor oscillations of relatively large amplitude exist during part of the sawtooth ramp, as in these situations a non-linear trigger condition would be more appropriate. Secondly, the sawtooth trigger condition says nothing about the non-linear consequences of the instability, and in particular on the relaxed q profile after the sawtooth crash. In existing codes, the Kadomtsev prescription [7] for the relaxed q profile, based on full reconnection of all magnetic surfaces inside the $q = 1$ volume, is often used. However, many experiments indicate that partial reconnection may also occur. A partial reconnection model was proposed in [6], which contains a free parameter, W_{crit}/r_1, where W_{crit} is a critical island width for the onset of widespread magnetic turbulence and r_1 is the $q = 1$ minor radius (full Kadomtsev reconnection corresponds to $W_{\mathrm{crit}}/r_1 \geqslant 2$). As this free parameter (W_{crit}/r_1), which is not determined within this theoretical model, is reduced, so the predicted sawtooth period becomes shorter.Thus a theoretical determination of W_{crit}/r_1 in realistic plasma regimes would certainly improve our degree of confidence in predicting the sawtooth behaviour in ITER. Consideration of three-dimensional space effects on sawtooth magnetic reconnection may be the way to resolve this long-standing theoretical issue.

There is another serious concern associated with the so-called trigger problem, namely, the non-linear development of the internal kink mode once the linear stability threshold is crossed. Experimentally, the mode is often observed to grow very rapidly as soon as it can be detected, in apparent contradiction with the slowly evolving equilibrium plasma parameters near the threshold and the fact that, near threshold, the linear growth rate is very small (= 0 exactly at threshold). One possible resolution of this problem is that diamagnetic effects, which provide an effective stabilization mechanism, are quenched as soon as the mode grows to a relatively small amplitude, such that the corresponding magnetic island becomes comparable with the ion Larmor radius [8]. Another possibility is that a secondary instability (such as a ballooning mode or a resistive interchange mode) is triggered when the magnetic island reaches a small width [9]. In both cases, the linear stability threshold proposed in [6] would still be fairly adequate in predicting the effective onset of the sawtooth crash, as these secondary effects (diamagnetic quenching, secondary instabilities) still require the internal kink mode to go linearly unstable before the crash can occur.

In conclusion, we may say that our level of theoretical understanding of the sawtooth phenomenon is certainly incomplete, but is probably sufficient to define a credible strategy for the prediction and the control of sawteeth in ITER. This statement is corroborated by our proven ability of controlling the sawtooth period and amplitude in existing experiments, as described in the next section.

2.1.2. Sawtooth control. Sawtooth control refers to the ability of an actuator system—here, some form of additional heating and/or current drive—to alter the period of the sawtooth instability. Since the ITER Physics Basis [1] there has been considerable experimental and theoretical progress on controlling sawteeth.

There are two main approaches to sawtooth control. In the first, one attempts to eliminate or avoid the sawtooth crash for the duration of the discharge by lengthening the sawtooth period (stabilization); while in the second, the goal is to increase the rate of the sawteeth (destabilization) to reduce the perturbation to the plasma at each individual crash.

In recent years, sawtooth control has been studied and used to advantage (e.g. see section 2.2) on many experiments. The experiments fall into two broad categories with either central or off-axis power deposition. All heating and/or current drive methods—NBI, ICRF and ECRF—acting alone can alter the sawtooth period (as can combinations of the different heating methods).

Sawtooth modelling (e.g. [6, 10–14]) is often carried out using (a) transport simulations (e.g. PRETOR or ASTRA codes) which evolve the electron, ion temperatures and current profiles between crashes, (b) sawtooth trigger criteria (e.g. [14] and references therein) and (c) profile relaxation rules at the crash—based for example on the Kadomtsev magnetic reconnection model. Partial reconnection is sometimes modelled but, at least for NBI stabilization [15], the general dependence of the period on the power remains the same as for full reconnection: the absolute sawtooth period is shorter with the former since the crash triggering condition is reached more quickly after the partial reconnection. The simulated sawtooth period is found to be in good agreement with many experiments in which the actuator parameters are varied *once the free parameter(s) of the model is(are) fixed*—for example, by matching the measured sawtooth period during

the purely Ohmic heating phase. Simulations can also separate the effects due to heating alone or current drive alone to elucidate complimentary or competing effects [14] which are especially important in the complicated ion cyclotron current drive (ICCD) case [16, 13].

In experiments, for localized off-axis heating, the deposition location is generally *swept* across the $q = 1$ region. In the simulations, it is difficult to predict the optimum actuator settings for stabilization with sufficient accuracy. This is due, first, to the combined uncertainties in the control parameters (e.g. launcher aiming, ray-tracing, etc) and equilibrium quantities (plasma position, q profile, etc) and second, to the strong sensitivity of the stabilization on the deposition location. However, while it may not be possible to predict exactly how the actuators must be set for stabilization prior to the experiments, it is important to note that simulations appear to be sufficiently advanced to accurately predict *whether a setting exists* for which the desired control is possible: the experimental deposition sweeps are employed to determine the optimum settings.

The sweeps are performed using the magnetic field, plasma position, or launcher mirrors—all sweeps are typically of small amplitude. The possibility of using mirrors is one of the main differences between ECRH and other heating/current drive methods (although antenna phasing might provide some flexibility during ICRF). Control of the launcher mirrors provides an *external* actuator in a feedback-control loop. For this reason we begin the discussion with EC sawtooth control then continue with NBI and ICRF.

ECRH/ECCD. During ECH, the change in temperature and thus conductivity leads to a change in local current density and magnetic shear ($s = (r/q)\mathrm{d}q/\mathrm{d}r$). With ECCD, both heating and current drive occur simultaneously. For deposition sweeps near the $q = 1$ surface both must be taken into account to properly simulate changes in shear and critical shear (for instability) and thus in the sawtooth period, e.g. [12]. The EC driven current is unipolar, and the current and power deposition profiles are generally nearly aligned to each other.

ECH just outside the $q = 1$ surface stabilizes sawteeth, illustrating the first method of sawtooth control (stabilization). The location of optimal destabilization with ECH is predicted by simulations and confirmed by experiments [14, 17], to be close to but inside the $q = 1$ surface. This demonstrates the second method of control.

Experiments show that even small amounts ($<1\%$ of the plasma current distributed over 10% of the minor radius in TCV) of co-ECCD just outside the $q = 1$ surface enhances the stabilization effect while counter-ECCD diminishes the effect [12, 18–21]. Simulations also predict that counter-ECCD close to but inside the $q = 1$ surface should stabilize sawteeth. This has been confirmed on ASDEX Upgrade [20, 22] and perhaps on JT-60U. On ASDEX Upgrade, moving counter-ECCD further inside $q = 1$ leads to complete elimination of the sawtooth crashes; however, from the changed MHD signature, it can be concluded that the stabilizing mechanism is different than that due to off-axis co-ECCD. On JT-60U, the sawtooth period matches well the expected $T_e^{3/2}$ scaling [23]: the sawtooth period is slightly longer than in the non-EC heated plasma. Co-ECCD just inside $q = 1$ is predicted [14] to be strongly destabilizing as observed on JT-60U [23]. (Detailed modelling would need to be carried out on ASDEX Upgrade and JT-60U to ensure that the predictions, which were made for TCV, are relevant in the experimental conditions of these machines.) Considerably stronger 'partial' stabilization outside $q = 1$ and destabilization inside the $q = 1$ surface, with current drive (compared with ECH aiming), has also been reported on DIII-D [24]. In all experiments and simulations, strong changes in sawtooth period are found for very small changes in deposition location (e.g. of the order of the EC deposition width consisting of a few centimetres) near the $q = 1$ surface.

Finally, simulations show that the stabilization process can be non-linear in power density; displaying the threshold behaviour observed in early stabilization experiments, e.g. [25]. Hanada *et al* showed experimentally that the threshold P_{EC}/P_Ω above which full stabilization occurs depends on the electron density and heating location (here P_{EC} = the applied ECRH power and P_Ω = the Ohmic power). The threshold was found to be lowest for deposition near $q = 1$ from the high field side (HFS) in low density discharges. It is actually the local shear that is changed so the relevant modification to consider is that of the local current density with heating power.

To compare sawtooth control between machines having vastly different baseline sawtooth periods without reference to simulations, the sawtooth period during the EC heated 'controlled' phase can be normalized to that in the non-EC heated 'non-controlled' phase: $\tau_{sN} \equiv \tau_{sEC}/\tau_s$ [17]. With fixed actuators and very careful placement of the deposition location the sawtooth instability can be avoided during the entire gyrotron pulse length in WT-3 (0.03 s) [25], TEXTOR (0.2 s) [26], JT-60U (1.5 s) [27] and ASDEX Upgrade (2.0 s) [28], corresponding to an increase in τ_{sN} to $\geqslant 75$, $\geqslant 11$, $\geqslant 6$ and $\geqslant 20$, respectively. Only WT-3 specifies that the power threshold has been reached in these experiments. In TCV with ECH deposition just outside the HFS $q = 1$ surface at power levels above the power threshold, the normalized sawtooth period reached values of $\tau_{sN} \sim 50$ ($\sim$110 ms) [29] corresponding to $\sim$50 electron energy confinement times and roughly one current redistribution time [14]. Here, however, a sawtooth crash occurred before the end of the 2 s gyrotron pulse, as the current profile evolved and the $q = 1$ location moved relative to the deposition location. Similar evolution effects may be at play in JT-60U in which a long period sawtooth is followed by sawtooth periods which scale as expected with $T_e^{3/2}$ [30]. These results indicate that, as in NTM stabilization experiments, feedback-controlled launcher angles and real-time calculations of the q profile are needed to ensure proper deposition for the entire pulse, if this method of sawtooth crash avoidance is to be used for ITER.

The success of the sawtooth period modelling on TCV, JET, FTU, etc gives confidence that differences between machines can be adequately taken into account in simulations to be able to give reliable predictions of the sawtooth control possibilities of the ITER EC actuators, provided that the sensitivity to the model free parameters is not too high in the ITER plasmas and that the heating and especially driven current are correctly calculated as inputs to the model. For example, WT-3 reports that the power threshold for optimum stabilization with low-field-side (LFS) deposition

was 1.7 times larger than on the HFS, and the ITER EC design does not allow for HFS $q = 1$ heating (where most present-day stabilization experiments are made). The difference in HFS and LFS stabilization was presumably due to electron trapping effects which can be taken into account by present-day ECCD modelling codes [31]. Note that it is crucial that if multiple EC beams are used, they must overlap precisely to benefit fully from the higher total power density.

Using the second control scenario (destabilization), τ_{sN} can be reduced to between 0.4 and 0.5 when heating (TEXTOR, ASDEX Upgrade, TCV, JT-60U) or co-ECCD (ASDEX Upgrade, TCV) is done inside but near the $q = 1$ radius. JT-60U has demonstrated a reduction to 1/6 the NB heated sawtooth period [23] with co-ECCD just inside $q = 1$. This method is likely to be the most robust (as well as the most efficient) destabilization [14]) EC sawtooth control scenario, as it is less sensitive to the deposition location.

NBI. NBI can also significantly affect sawtooth behaviour. Trapped and/or co-circulating-passing NBI-generated fast particles are expected to provide the stabilization, as these populations typically provide an energy sink, or additional potential well. On JET, dedicated experiments were carried out to investigate the sawtooth period increase during central NBI heating [32]. The discharges were in the weakly collisional ion-kinetic regime and the time of the crash at high NBI power was determined by the condition $-c_\rho \hat{\rho} < -\delta \hat{W}$ (see section 2.1.3 for details on this notation). The same scenario is considered most likely in ITER [6]. Detailed sawtooth period simulations [16] provide a good match to the experimental sawtooth periods. The fast particle contribution to δW is inversely proportional to s_1 (the shear at $q = 1$), so methods to alter the evolution of s_1, in order to alter the sawtooth period, as described for ECH, can also be used in the case of fast particle stabilized sawteeth (see ICCD section); [33].

Both the current profile and/or the pressure profile can be modified by NBI as noted by JT-60U and ASDEX Upgrade. Negative ion NBI in JT-60U produces a fast co-passing population that extends the sawtooth period [34–36] and has been shown theoretically to be stabilizing [37]; whereas, counter-circulating ions are predicted to be destabilizing. ASDEX Upgrade found that the sawtooth period varied as a function of the different individual NB injectors being used (8 are available). In particular, each of the two tangential beams deposit fast ions at about midradius and, except near the plasma edge, these ions are co-circulating. These beams are found to be stabilizing [28]. Combining several beams, the sawtooth period is not only increased but fishbones are eliminated and are largely replaced by an $m = n = 1$ mode. The mode has two frequencies and two radial locations suggesting the existence of two $q = 1$ surfaces in the plasma [38] and consequently a modification of the current profile. Recent theoretical work involving asymmetric circulating ions suggests that on-axis counter-injection, or off-axis co-injection, of neutral beams can destabilise sawteeth, and thus offset the stabilising properties of trapped alpha particles in ITER [37].

ICRH/ICCD. In the ion cyclotron range of frequencies (ICRF), resonant heating (ICRH) has been used to control the sawtooth period on JET and JT-60U. ICRH has been used [4,39] to lengthen (shorten) the sawtooth period by heating near the $q = 1$ surface with an antenna phasing of $+90^\circ$ (-90°); taking advantage of the ICRF-induced pinch of the resonating trapped ions to increase (decrease) the fast particle pressure inside the $q = 1$ surface and thereby alter the stabilization of the sawtooth instability [40].

Recently, the sawtooth period has been successfully controlled in JET using the ITER-relevant 2nd *harmonic* minority ion frequency scheme. As in many EC experiments, a sweep of the ICRH deposition near the $q = 1$ surface was made varying the toroidal field. Both $\pm 90^\circ$ phasing and both HFS and LFS locations were studied during a minority density scan. An NBI ramp was performed at fixed toroidal field and plasma current leading to additional effects for HFS deposition as the NBI deuterons are resonant with the ICRF waves near the plasma centre [13]. With ICCD alone, the driven current (protons) is of the diamagnetic type caused by the finite orbit widths of trapped resonating ions and results in a bipolar current perturbation as well as collisional electron heating. Two maxima in the sawtooth period are observed as the resonance location passes through the LFS $q = 1$ inversion radius, with a further maximum as the resonance approaches the plasma centre [41]. Simulations [16] indicate that, first, electron heating well outside $q = 1$ stabilizes the sawteeth, next, a more equal distribution of power near the $q = 1$ surface reduces the stabilization and, finally, the combined effects of co- (counter-) ICCD outside (inside) $q = 1$ are together stabilizing enough to again increase the sawtooth period: that is, modifications in local shear dominate in the crash trigger as in the case of ECH/CD. As the resonance moves further towards the magnetic axis, the increase in fast ion pressure, which modifies δW, appears to be the stabilization mechanism, as with NBI. Good sawtooth control is found on the LFS, nearly independent of antenna phasing ($\pm 90^\circ$; though still due to ICCD from trapped resonating ions) by tuning the resonance location. This results in a modification of the local shear at $q = 1$ which can be used to delay the triggering of the 3/2 NTM mode during power ramp-up experiments.

In JT-60U detailed measurements of Alfvén eigenmode frequencies during second harmonic minority heating indicate that $q(0) < 1$ during the entire sawtooth period. An alternative full reconnection model of Kolesnichenko *et al* [42] has been used to account for this observation. This model has *not* yet been implemented in a code to track the evolution of the sawtooth period and to study the control of sawteeth.

Most sawtooth period modelling uses a 1D transport model; details of the sawtooth crash process are lost. Additional 2D effects, observed in some experiments, such as stable island rotation during tangential NBI or near central ECH/ECCD are also not possible with the 1D modelling. Recent simulations [43 and references therein] recover many detailed 2D effects seen in experiments, which can be due to incomplete reconnection, by introducing an ad hoc $m = n = 1$ island displacement function; however, at present no model exists to predict the evolution of the island displacement. Nevertheless, the success of simple 1D models in predicting sawtooth control in present-day experiments, gives confidence that sufficient physics is contained in the model for predictions of the control capabilities of actuators in the ITER device.

2.1.3. Central MHD activity expected in ITER. As discussed in the preceding section, there are now improved tools for modelling sawteeth, which have been better benchmarked with existing experiments. These have been used to evaluate the effects of sawteeth in the present ITER design.

A comprehensive model of the sawtooth trigger and relaxation oscillation has been developed by Porcelli *et al* [6]. We adopt the notation used in that paper and examine the consequences of this model for ITER. An effective potential energy functional, associated with the macroscopic drive, can be defined by

$$\delta W = \delta W_{\rm core} + \delta W_{\rm fast}. \tag{1}$$

Here $\delta W_{\rm core} = \delta W_{\rm MHD} + \delta W_{\rm KO}$, $\delta W_{\rm MHD}$ is the ideal MHD term containing destabilizing terms due both to plasma pressure and parallel current density [44], $\delta W_{\rm KO}$ is the 'Kruskal–Oberman' term [45] which represents the (normally stabilizing) effect of collisionless thermal trapped ions, and $\delta W_{\rm fast}$ represents the kinetic effects related to the high energy fusion-produced alpha particles.

Reference [6] gives three conditions, any of which means triggering the sawtooth crash if it is met. These conditions, which we call the Porcelli criteria, are

$$\begin{aligned} &-\delta\hat{W}_{\rm core} > c_{\rm h}\omega_{\rm Dh}\tau_{\rm A},\\ &-\delta\hat{W} > 0.5\omega_{*\rm i}\tau_{\rm A},\\ &-c_\rho\hat{\rho} < -\delta\hat{W} < 0.5\omega_{*\rm i}\tau_{\rm A} \text{ and } \omega_{*\rm i} < c_*\gamma_\rho. \end{aligned} \tag{2}$$

Here $c_{\rm h}$ and c_ρ and c_* are numerical factors of order unity, $\hat{\rho} = \rho_{\rm i}/\bar{r}_1$ with $\bar{r}_1$ being the average radius of the $q = 1$ surface, $\rho_{\rm i} = v_{\rm thi}/\Omega_{\rm ci}$ is the thermal ion Larmor radius, $\omega_{*\rm i}$ is the ion diamagnetic frequency, γ_ρ is the characteristic growth rate of the internal kink mode in the ion-kinetic regime, $\omega_{\rm Dh}$ is the precessional drift frequency of the high energy particles, and the potential energy is normalized here according to

$$\delta\hat{W} \equiv -\frac{4\delta W}{s_1\xi^2\varepsilon_1^2 RB^2} \tag{3}$$

with s being the magnetic shear, ξ the radial displacement of the magnetic axis, ε the inverse aspect ratio and the subscript 1 denotes values at the $q = 1$ surface.

The first of these criteria (2) is that the high-energy trapped particles complete many precessional drift orbits within a characteristic magnetic perturbation time, which is of order $|\delta\hat{W}_{\rm core}^{-1}|\tau_{\rm A}$, with $\tau_{\rm A}$ being the Alfvén time. The second is that the diamagnetic rotation is not sufficient to stabilize the mode. The third criterion is for the case when the energy drive is sufficiently weak that the mode is stable according to the first two criteria. In this case, the $m = 1$ mode structure changes its nature from a global internal kink to a drift-tearing mode localized near the $q = 1$ surface. This is normally stable because of kinetic layer effects, but the layer effects will be insufficient when the normalized potential energy exceeds the normalized ion Larmor radius and rotation effects are sufficiently weak (note that the choice of normalizations has made this comparison appropriate).

When the sawtooth is predicted to be triggered by the Porcelli criteria, the transport coefficients are modified in two ways. Assuming that the central safety factor q_0 is less than unity, the value of the toroidal flux at the inversion surface, Φ_1, is calculated from the integral equation:

$$\int_0^{\Phi_1}\left(\frac{1}{q(\Phi)} - 1\right){\rm d}\Phi = 0. \tag{4}$$

For the duration of the sawtooth crash time $\tau_{\rm CRASH}$, the thermal conductivity and the hyper-resistivity in the central region $\Phi < \Phi_1$ are defined to be: $\chi = r_1^2/\tau_{\rm CRASH}$ and $\lambda = \lambda_0 B_0^2 r_1^4/\tau_{\rm CRASH}$. A value of $\lambda_0 = 0.1$ effectively causes a Kadomtsev reconnection to occur [46] in the time $t = \tau_{\rm CRASH}$, which is taken to be 10 ms. This value was chosen as a compromise to minimize numerical time-resolution requirements. Even though the actual experimental value may be considerably less than this, we do not expect this choice to affect the global result since it is much less than the energy confinement time, $\tau_{\rm CRASH} \ll \tau_E$. By lowering λ_0 to 0.001, an incomplete reconnection can be modelled where the temperature profile flattens but the current and flux do not fully reconnect.

It has previously been reported [15] that the Porcelli sawtooth model described here has been implemented in the PRETOR code and compared in detail with JET experiments in several different regimes with different levels of NBI power. It was reported that in every case analysed, the simulated sawtooth periods are within 20% of the experimentally observed periods, even as the period varies by more than a factor of 5 during a given shot [15]. This same Porcelli sawtooth model has been implemented in the Tokamak Simulation Code (TSC) [47] and its consequences on transport and ignition in ITER have been investigated. In order to predict the profiles and performance, two of the leading H-mode models are implemented in TSC: (A) the multi-mode model [48], and (B) the Gyro-Landau fluid model GLF23 [49]. These core transport models must be supplemented by boundary and edge models.

The H-mode models (A) and (B) are only applied in the central region $0 < \Phi < 0.75$, where Φ is the normalized toroidal magnetic flux that is zero at the magnetic axis and unity at the plasma/vacuum separatrix. In the edge region $0.75 < \Phi < 1.0$, a transport model $\chi_I = \chi_{\rm e} = C/n_{\rm e}$ is used, where $n_{\rm e}$ is the local electron density and C is a constant chosen so as to make the pressure gradient in this region below the infinite-n ballooning mode stability criteria. This leads to electron and ion temperatures at the top of the pedestal, $\Phi = 0.75$, of 3–5 KeV. The density profile is not advanced in time in these simulations, but is rather a prescribed function of normalized poloidal flux, ψ, and time, t. The electron density during the current flattop is taken to be

$$n_{\rm e}(\psi, t) = n_0(t) \times [(1 - \psi^\beta)^\alpha + r_{\rm edge}], \tag{5}$$

with parameters $\alpha = 0.25$, $\beta = 8.00$, $n_0 = 0.75 \times 10^{20}\,{\rm m}^{-3}$ and $r_{\rm edge} = 0.4$ chosen to correspond to a reference ITER ELMy H-mode case.

The Gyro-Landau fluid model transport GLF23 is used. A feedback system on the ICRH power designed to keep the total stored energy W constant at 320 MJ is utilized. A uniform distribution of 2% beryllium and 0.12% argon impurities is included, which together with the He build-up (assuming $\tau_{\rm P} = 18.5$ s), leads to a value of $Z_{\rm EFF} \sim 1.65$ during the flattop. The constant in the edge region is chosen as $C = 2.5 \times 10^{19}$.

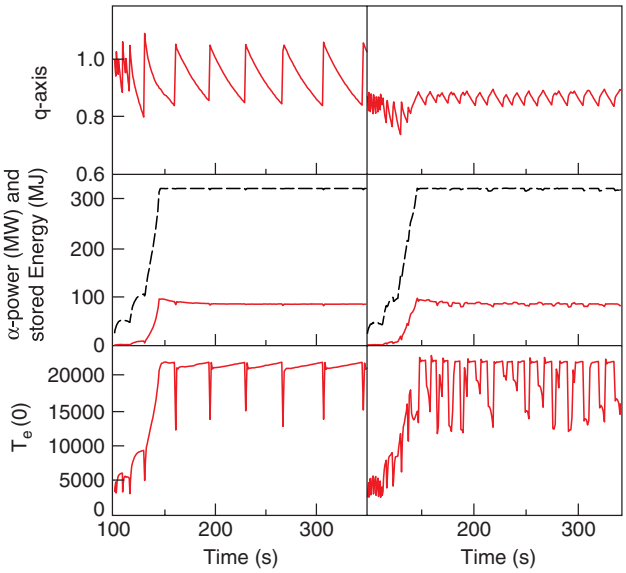

Figure 1. ITER simulations using the GLF23 transport model and both the complete (left) and incomplete (right) reconnection models. The top frames show the axial q value, in the middle frames the α-power is the solid line and stored energy the broken line and the bottom frames show the axial temperature.

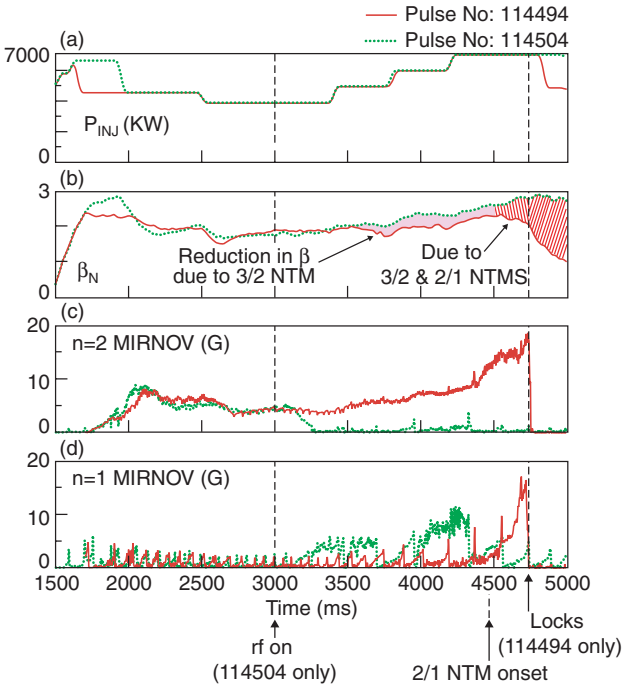

Figure 2. DIII–D discharges with (114504, dotted lines) and without (114494, solid lines) ECCD suppression of an $m/n = 3/2$ neoclassical tearing mode. (*a*) Neutral beam power, (*b*) β_N, (*c*) $n = 2$ Mirnov $|\tilde{B}_\theta|$, (*d*) $n = 1$ Mirnov $|\tilde{B}_\theta|$. The degradation in energy confinement due to the NTM from 3/2 and 2/1 NTMs can be seen in the effect on β_N.

The results of two simulations are shown in figure 1, one using complete reconnection and the other with incomplete reconnection. Note that the sawtooth period is about 50 s during the flattop for the complete reconnection, and 2–3 times that frequent for the incomplete reconnection. Both the alpha power and stored energy are essentially independent of the sawtooth period, since the period is always longer than the energy confinement time, the core electron temperature is so high, and the magnetic diffusion time is long compared with the energy confinement time. The $q = 1$ radius is about 42% of the minor radius.

It is concluded that in ITER the sawtooth will lead to periodic oscillations on a time that is considerably longer than the energy confinement time, $\tau_{SAW} \gg \tau_E$, and that the temperature at the $q = 1$ surface is sufficiently high that the sawteeth oscillations have negligible effect on both the stored energy and the rate of neutron production (as can be seen from figure 1). The incomplete reconnection sawteeth have a period about half that of the complete reconnection, but still long compared with τ_E. A consequence of these long period sawteeth is the possibility of destabilizing NTMs (see section 2.2.3) and thus sawtooth control (see section 2.1.2) is important to consider. Also, heat pulses from sawteeth can be linked to ELM triggering.

2.2. *Neoclassical tearing modes*

2.2.1. Physics of neoclassical tearing modes. The presence of persistent magnetic islands in the plasma core is an important issue for burning plasmas, as they can significantly limit the performance of both the standard ELMy H-mode and advanced scenarios. The neoclassical tearing modes (NTMs) are driven by the local reduction of the bootstrap current due to the pressure flattening across the island. This drive is inherently a non-linear process as it relies on the existence of a fully formed island, large enough to increase the local radial transport and flatten the pressure profile. The most significant NTMs are those with $m/n = 3/2$ or 2/1 (with m the poloidal mode number and n the toroidal mode number). The effect of these NTMs on energy confinement is nicely illustrated by comparison of two discharges one of which suffers from a 3/2 and then a 2/1 NTM, and an otherwise identical discharge in which electron cyclotron current drive (ECCD) stabilizes the 3/2 NTM and a 2/1 NTM does not occur (see section 2.2.2 for a discussion of NTM stabilization), as shown in figure 2.

Since the previous report [1], the studies have concentrated either on the capabilities to predict the onset of NTMs in ITER and on the possibilities to stabilize the modes if they are triggered. Scalings at the mode onset and decay including both collisionality and Larmor radius have been extensively investigated, e.g. in ASDEX Upgrade [50, 51], DIII-D [52], JET [53,54], JT-60U [55] and T-10 [56], leading to a consensus on the stronger, approximately linear, dependence on Larmor radius compared with collisionality. Cross-machine comparisons of onset conditions have led to a scaling for the onset beta of the $m = 3/n = 2$ NTM depending on ν_{i*} and mainly ρ_{i*} [57]. The difficulty of such a scaling is that it needs to combine the seed island formation physics and the NTM physics. A particular assumption made in [57] is that the seed island width can be described as a function of β_p and $1/S = \tau_A/\tau_R$, the inverse of the magnetic Reynolds number. Recent JET experiments have shown however that large seed islands and hence NTMs, can be triggered at the sawtooth crash after long sawtooth free periods [4], even at low β. Therefore more recent studies have concentrated on

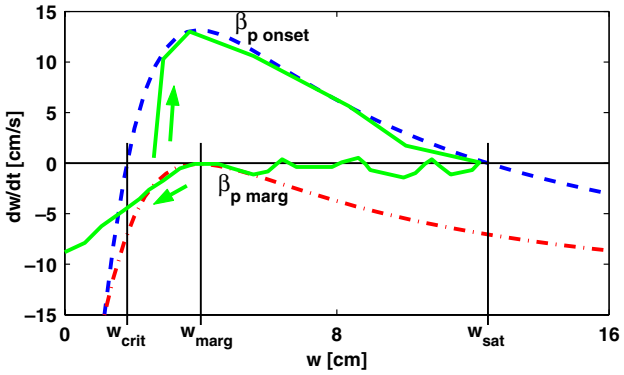

Figure 3. Sketch of the time evolution of the island growth rate as given by equation (6) at the onset of the NTM when the critical seed island size (W_{crit}) is exceeded and an NTM forms at $\beta_{p,onset}$. A slow decrease in beta from $\beta_{p,onset}$ to $\beta_{p,marg}$ (when $\max(dW/dt) = 0$) is assumed, as in power ramp-down experiments, such that $dW/dt \approx 0$ (reproduced from [54] 'Marginal β-limit for neoclassical tearing modes in JET H-mode discharges').

the seed island formation or on the NTM physics, as they are not necessarily related. However, additional effects on NTM onset arise from resonant error fields which can seed NTMs and slow plasma rotation [58]. Preliminary analysis suggests decreased plasma rotation could be reducing the small island polarization current threshold and thus making the 2/1 NTM unstable at lower β, but further experiments are required to elucidate this effect.

To understand the physics mechanisms at play, it is best to describe in some detail the modified Rutherford equation, which can be written symbolically as follows for the island growth rate:

$$\frac{\tau_R}{r_s}\frac{dW}{dt} = r_s\Delta'(W) + r_s\beta_p(\Delta'_{BS} - \Delta'_{GGJ} - \Delta'_{pol}) + r_s\Delta'_{CD}. \quad (6)$$

Here W is the width of a magnetic island occurring at a radius r_s and τ_R is the local resistive diffusion time; Δ' is the stability index of the equilibrium current profile, Δ'_{BS} is the bootstrap drive term, and Δ'_{CGJ} and Δ'_{pol} are the stabilizing curvature [59] and polarization terms [60], respectively. The effect of current drive represented by Δ'_{CD} will be discussed in the next section. The island width dependence is $\Delta'_{BS} \propto W/(W^2 + W_d^2)$ and $\Delta'_{pol} \propto W_{pol}^2/W^3$, where W_d describes a stabilizing effect at small island width due to perpendicular thermal conduction [61] and W_{pol} is a constant related to the stabilizing polarization effect. A fuller description of these terms, used to compare with experimental data, can be found in [54,62] and references therein (see also [1]). The typical evolution of the island growth rate in a full discharge, assuming a slow ramp-down of the power, and thus a slow decrease in the terms proportional to β_p in equation (5), is shown in figure 3. At a given time in the discharge, an island is triggered at a beta value $\beta_{p,onset}$, in most cases much larger than $\beta_{p,marg}$, and subsequently grows to a relatively large saturated island width. When $\beta_p \leqslant \beta_{p,marg}$, the mode is stabilized and the growth rate becomes rapidly very negative. The hysteresis, ratio $\beta_{p,onset}/\beta_{p,marg}$, is significant in standard scenarios with modest size sawteeth, it has been measured much above unity in ASDEX Upgrade [51], DIII-D [52], JET [4] and JT-60U [55]. This occurs because $\beta_{p,marg}$, the marginal beta limit above which NTMs are metastable, is very low but generally the sawteeth (or other seeds) do not form a large enough island ($W < W_{crit}$) until β_p increases well above $\beta_{p,marg}$. Since $\beta_{p,marg}$ scales approximately linearly with ρ^* [54] ITER is predicted to have $\beta_p > \beta_{p,marg}$ as soon as it is in the H-mode. Therefore the existence of NTMs in ITER does not depend on β as such, but rather on the triggering of a seed island $W_{seed} > W_{crit}$. Thus, the predictions of seed island widths and of the value of W_{marg} are each of great importance for burning plasmas.

The prediction of W_{marg} indicates that its size normalized by the minor radius will be much smaller in ITER than in present experiments. Its value depends on all the terms in equation (6) and their dependence at small W. The understanding and relevance of each of these terms have been further developed since [1]:

- The first term is the classical Δ' term, which has a weak dependence on W. It has been shown that classical tearing modes can provide the seed islands for NTMs [63,64], and this may be one of the possible explanations of the 'triggerless NTMs' observed in other machines like ASDEX Upgrade [65], JT-60U [55], T-10 [56] and TFTR [66]. This usually happens when the current profile is modified [63], and could become the main seed island trigger mechanism in hybrid scenarios, or when β approaches the ideal limit, as Δ' can become large and positive [64] and thus could become important in advanced scenarios. In addition, using fast power shut-off, leading to a rapid vanishing of the terms proportional to β_p in equation (6), it was possible to show the linear $\Delta'(W)$ on W in TCV [17].
- The second term is the bootstrap drive, which is reduced at small island width due to two main effects. First, the ratio of perpendicular to parallel heating becomes non-negligible and the pressure profile is not flattened completely, reducing the perturbed bootstrap current [61, 67]. Anomalous perpendicular viscosity can also affect the bootstrap drive. Its effect is frequency dependent and can be stabilizing or destabilizing depending on the sign of ω/ω_{*pi} [68], where ω is the mode frequency in the electron frame and ω_{*pi} is the ion diamagnetic frequency. Another effect which reduces the perturbed ion bootstrap current even more is finite ion Larmor radius effects [69]. When the island width is less than $\sim 5\rho_b$ (ρ_b = ion banana width) ions are still affected by the pressure gradient inside and outside the island, leading to a finite bootstrap current within the island.
- The third term describes the stabilization due to the effect of curvature and is usually smaller than the bootstrap term in present tokamak scenarios and therefore has often been neglected in the past. It has been confirmed in MAST to be significant for tight aspect ratio scenarios [70]. On the other hand, it has been shown to yield a finite stabilizing term for small island width [71] and therefore can be significant at small island widths in present tokamaks and for ITER aspect ratio as well.
- The fourth term is due to the polarization current, resulting from the fluctuating electric field driven by the different electron and ion responses to the rotating island. Therefore it involves diamagnetic effects, effective mode

frequency and ion and electron kinetic effects as the island width of interest is of the order of the ion orbit width. Excitation of drift waves has recently been included as well [72] and the effects of the ion sound [73]. It has also been noted that in toroidal geometry, the variations of plasma pressure within the perturbed magnetic surfaces (due to ion sound waves and/or finite parallel heat conductivity) lead to a new type of frequency dependent polarization terms that are important for determining the threshold value of the seed island [74]. The complexity of determining this contribution has been highlighted in the recent review [62]. The polarization current term can change sign, depending on the effective mode frequency, which is very difficult to measure. It is observed to be in the ion drift direction in most cases in DIII-D [75] and JT-60U [76]. A common feature however of all the different models, is that the polarization current rapidly decreases with increasing island size and it scales with $\rho*$. Therefore the polarization current effect could be less important in burning plasmas if W_{marg} is of the order of a few poloidal Larmor radius. However present understanding of experimental results require a better quantitative theoretical knowledge of this term. For example it has been shown recently that finite ion orbit widths also strongly affect this term [77].

The prediction of the seed island formation is much more complicated, also because it can have different origins for different plasma scenarios like non-linear mode coupling [78], link to the current profile of 1/1 activity [53], fishbones [65], turbulence [79], error fields [58, 80] in addition to the causes mentioned above. Further changes in the rotation profile between the location of the source of the seed island and the NTM resonant surface can also affect the onset β—experiments in JET suggest that decreased rotation reduces the NTM onset β-limit [81]. For ITER, which will have sawteeth stabilized by α-particles, the main perturbation will most probably be at the sawtooth crash, as is the case in high field/current scenarios in JET [4]. The seed island can be formed by the direct non-linear coupling to the 1/1 mode [82]. It should be mentioned that 2/1 modes can be directly triggered at the sawtooth crash at low beta in JET [83]. Therefore better understanding is clearly needed in this domain, together with sawtooth control experiments (see section 2.1 and 2.2.3).

The effects of increased shaping, elongation and triangularity, are in general weak for NTMs but they can affect different terms. Therefore they can manifest themselves differently in local and global parameters [84]. Since the global confinement properties are modified, both the current and pressure profiles are affected and therefore also the $\Delta'(W)$ and Δ'_{BS} terms. Shaping also directly modifies the curvature term, Δ'_{GGJ}, but it is expected to be significant only for tight aspect ratio tokamaks in conjunction with small magnetic shear [85, 86]. Both elongation and triangularity can also be important in the seed island trigger mechanisms, as they influence the sawtooth stability (see section 2.1) and the coupling between the sawteeth and the m/n flux surface.

The power ramp-down experiments, as illustrated in figure 3, are important since they can provide accurate measurements with NTM island widths of the order of W_{marg}, where all the stabilizing terms are important. Therefore cross-machine scalings are possible, in particular as these measurements of W_{marg} do not depend on the seed island formation mechanism. Such studies have been started [58] on ASDEX Upgrade [51], DIII-D, JET [54] and JT-60U, but much more work is required in order to extend the database to a larger range in $\rho*$ and collisionalities. Another domain for which further physics understanding should be developed is the effects of fast particles on NTMs and their interaction with large islands [87]. In summary, the NTMs are expected to be metastable in ITER, i.e. $\beta > \beta_{\mathrm{marg}}$, however we are unable to predict with certainty if large enough seed islands will be triggered and destabilize these modes. Therefore it remains prudent to have a means to stabilize them as discussed in sections 2.2.2–2.2.4.

2.2.2. *Active control of NTMs.*

Status at time of IPB. In the ITER Physics Basis, the possibility of NTM stabilization was only briefly described in the ECCD chapter. No experimental results were available at that time. This subsection describes direct stabilization of the NTM. Other methods that work on stability of other modes that influence the NTM (sawteeth to suppress seed island and $(m+1, n+1)$ to trigger FIR regime) are discussed in section 2.2.3. Here, theory and experimental verification as well as elements of a feedback-control scheme are considered. The extrapolation to ITER is given in section 2.2.4.

Theory. NTMs may be stabilized by local CD at the resonant surface $q = m/n$ due to two effects: The equilibrium (i.e. (0,0)) component of the driven current alters Δ', while the (m, n) component and its multiples directly counteract the reduction of bootstrap current within the island due to the flattening of the pressure gradient there. This has been studied with a variety of theoretical approaches including non-linear resistive MHD codes or analysis based on the Rutherford equation with a CD-term (see e.g. [88]). The basic physics results are in agreement, so that here, for the sake of simplicity, the Rutherford equation is discussed (equation (6) of section 2.2.1). The CD-term in this equation can be written as

$$\Delta'_{\mathrm{CD}} = -L_q \left(a_{m,n} \eta_{\mathrm{CD}} \frac{I_{\mathrm{CD}(m,n)}}{I_{\mathrm{p}}(r_{\mathrm{s}})} \frac{1}{W^2} + a_{0,0} \frac{I_{\mathrm{CD}(0,0)}}{I_{\mathrm{p}}(r_{\mathrm{s}})} \frac{1}{\delta_{\mathrm{ec}}^2} \right), \tag{7}$$

where $I_{\mathrm{CD}(m,n)}$ is the total helical current of same helicity as the island, $I_{\mathrm{CD}(0,0)}$ is the (0,0) component of the externally driven current changing the equilibrium stability index Δ' in equation (6) of section 2.2.1, and $L_q = q/(\mathrm{d}q/\mathrm{d}r)$. The term with η_{CD} takes into account the conversion of the externally driven current into a helical current as the current generated locally rapidly distributes on the island flux surfaces. For a δ-function in the O-point (i.e. the extreme limit of modulated ECCD), $\eta_{\mathrm{CD}} = 1$. Due to the redistribution of current on the island flux surfaces, a source that is constant and does not depend on the helical angle may still create a helical current if it is deposited in a radial region smaller than the island width. Conversely, if the deposition width δ_{ec} is much bigger than W, a uniform source only creates a constant current that does not have a helical component. However, if the power is modulated, a helical current is created even if $\delta_{\mathrm{ec}} \gg W$. This

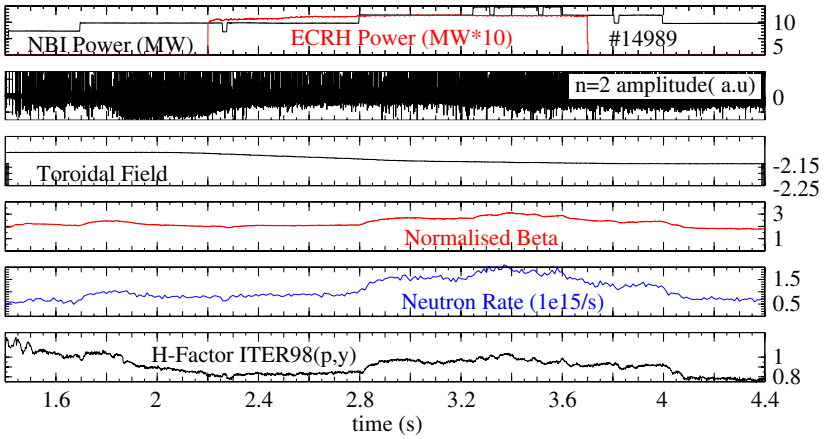

Figure 4. Complete stabilization of a (3,2) NTM in ASDEX Upgrade. β_N can be increased well above the onset level without reappearance of the mode. At even higher β_N, the mode comes back due to a deposition mismatch induced by the Shafranov shift.

effect can usually not be seen in present-day experiments since the width $W_{\mathrm{marg,CD}}$ at which $\mathrm{d}W/\mathrm{d}t$ becomes negative during stabilization is comparable to δ_{ec}. In ITER, however, this may be different and it may be required to modulate the ECCD power in phase with the island O-point for complete mode stabilization. The requirement for good localization leads to the figure of merit $(I_{\mathrm{CD}}/I_{\mathrm{p}}(r_{\mathrm{s}}))(r_{\mathrm{s}}L_{\mathrm{q}}/\delta_{\mathrm{ec}}^2)$ for equilibrium current profile modification as well as for helical current generation if $\delta_{\mathrm{ec}} > W_{\mathrm{marg,CD}}$ and the power is injected continuously. For helical current generation in the case $\delta_{\mathrm{ec}} > W_{\mathrm{marg,CD}}$, but with modulation, $(I_{\mathrm{CD}}/I_{\mathrm{p}}(r_{\mathrm{s}}))(L_{\mathrm{q}}/\delta_{\mathrm{ec}})$ is the relevant figure of merit. Since the current localization is only good if the current is deposited in the island, the source also has to be localized precisely in radial direction. Theory suggests that if the mismatch is of the order of the island width, stabilization becomes ineffective [89].

One goal of the theoretical description is the prediction of the power level required for complete stabilization of a given mode. Here, the main uncertainties lie in the terms describing the mode dynamics rather than in the CD part, which seems to agree quite well with the experiment. Thus, we have to improve our predictive capability for the NTM stability itself to improve the predictions of the required power for stabilization as given in section 2.2.4.

Experimental results. Due to the requirements for localization, ECCD is clearly the best method for NTM control, but LHCD has also been used to modify the global current profile, thus altering NTM stability mainly through a change in Δ' [90]. In the experiments, a local current is generated by the waves, but also the local heating within the island may lead to a local change in resistivity and therefore generate an additional helical current.

Successful ECCD experiments have been carried out in ASDEX Upgrade [91], JT-60U [92] and DIII-D [93, 94] demonstrating complete stabilization of both the (3,2) and the (2,1) NTM. The required power was of the order of 10–20% of the total heating power, with the (2,1) stabilization requiring more power than the (3,2) stabilization. The NTM remained stabilized even in the presence of sawteeth and fishbones, which provide seed islands, and β could be increased by 20–30% above the onset beta without reappearance of the mode. Usually, if ECCD is switched off, the next sawtooth or fishbone usually triggers the mode again. Figure 4 shows an example from ASDEX Upgrade (an example from DIII-D is shown in figure 2).

In the experiment shown in figure 4, the mode ultimately comes back due to the Shafranov shift at higher β which moves the resonant surface away from the ECCD deposition. This clearly calls for a feedback control of the deposition (see below). In some experiments, β did not recover to the previous value and NBI power had to be increased to recover the β. This is due to a confinement reduction when using ECCD that is a generic feature of strong electron heating in scenarios with $T_{\mathrm{i}} > T_{\mathrm{e}}$ (see chapter 2 of this issue [95]).

Several experiments have proven the need for exact localization of the ECCD by sweeping ECCD over the resonant surface using a slow B_{t}-ramp or antenna mirror steering. It is found that ECCD is only effective in reducing the mode amplitude when it is deposited close to the resonant surface; if deposition is outside the island, no stabilization occurs [88]. Experiments were also performed on ASDEX Upgrade with variation of the toroidal injection angle. For perpendicular injection, only a weak stabilizing effect is found at a power level that leads to full stabilization for optimum toroidal angle [96]. This proves the important role of the direct CD as opposed to the current generated by local heating. As pointed out above, theory predicts that $I_{\mathrm{CD}}/\delta_{\mathrm{ec}}$ or $I_{\mathrm{CD}}/\delta_{\mathrm{ec}}^2$ is the figure of merit for stabilization. This means that although the driven current increases monotonically with toroidal injection angle ϕ, the largest toroidal injection angle is not necessarily the best for stabilization, because δ_{ec} also increases with it. Usually, I_{ECCD} increases more weakly than linearly with ϕ whereas δ_{ec} increases more strongly than linear, so that an optimum angle exists. This was shown experimentally on ASDEX Upgrade where above a certain toroidal injection angle, the mode could no longer be completely stabilized at the power level sufficient at smaller angle in spite of the higher driven current at larger ϕ [97].

Experimentally, it is difficult to sort out the stabilizing contributions due to the (0,0) component of the EC current and due to the helical component of the EC current since both cannot directly be measured. The fact that the localization within the island is required to obtain good stabilization

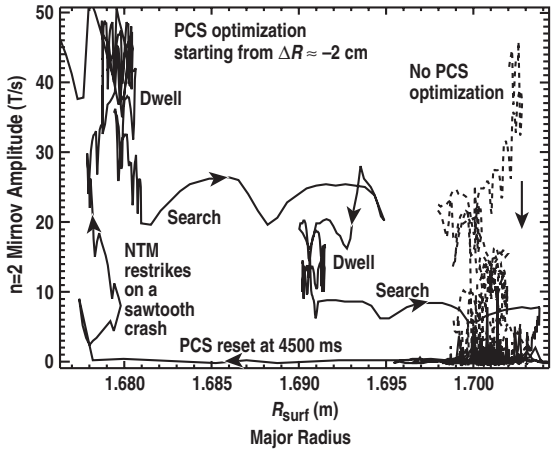

Figure 5. Search and suppress algorithm in DIII-D (Reprinted with permission from [93]. Copyright 2002, American Institute of Physics): The major radius is adjusted in steps by the plasma control system (PCS) until complete stabilization occurs. After reset of the position, the mode reappears at a sawtooth crash. A discharge with no PCS optimization, but with ECCD applied at the optimum major radius found by the PCS, shows complete suppression

efficiency points towards the importance of the helical current for the ECCD schemes. A hint towards the importance of Δ' is the observation that in scans of the deposition across the resonant surface, a stabilizing effect is found for counter-ECCD deposited at a radius significantly smaller than the resonant surface [98]: while the helical current produced by counter-ECCD would be destabilizing, the (0,0) component of the ECCD current in this case gives a stabilizing contribution to Δ'. Modelling of ECCD experiments suggests that the two effects may be of same order in ASDEX Upgrade. In LHCD experiments on COMPASS-D, the deposition width was so large compared with the island size that Δ' changes were the only reasonable explanation for the observed stabilization [90]. The importance of both contributions is of special interest for the extrapolation to ITER, where usually, the Δ' term is not yet taken into account (see, e.g. the discussion in section 2.2.4).

Elements of a control scheme. The experimental results discussed so far show that local ECCD is a good candidate for an NTM control scheme in ITER. However, the high sensitivity to the correct localization calls for a feedback-control scheme of the deposition. This is required when changes in the equilibrium, e.g. due to the Shafranov shift variation with β, move the resonant surface with respect to the deposition.

Several experiments have explored the possibilities for such a control. In DIII-D, a so-called 'search and suppress' algorithm was used [93] (see figure 5): the deposition radius was varied by a small step through either adjusting B_t or R_0. Based on the effect of this step on the mode, it was decided how the next step should be taken. In JT-60U, the mode position was determined online from ECE measurements from the local minimum of mode amplitude (which occurs at the rational surface) and the poloidal steering angle of the ECCD launcher was adjusted to track the mode [23, 99]. More sophisticated ideas on mode detection have been proposed and in some cases tested, e.g. on FTU using ECE detection [100, 101]. More experimental research is needed to clarify the applicability of feedback-controlled schemes to ITER.

Another element of a future NTM control scheme is the early detection of the mode: the power needed to reduce the mode is predicted to be small at mode onset. Figure 6 shows an experiment in JT-60U where ECCD was applied before mode onset [102,103]. In this situation, the mode only grows to small amplitude after triggering, whereas the same power applied to a saturated mode could not reduce the mode amplitude to such small values.

In summary, local ECCD has successfully been used to suppress NTMs and recover β-values at modest ECCD power level. With ECCD, stable operation without NTMs at β above the onset β-value without ECCD, has been demonstrated. Although the qualitative physics picture of the process is confirmed by the experiments, better theory and more detailed comparison between theory and experiment is needed to better predict the power requirements for ITER in a quantitative manner.

2.2.3. *Mitigation of NTMs.*

Seed island control. The neoclassical tearing modes (NTMs) have a relatively low marginal β and are therefore metastable, typically as soon as the discharge is in the H-mode, with $\beta_{N,marg} \geqslant 0.5$–1, for standard $q_{95} \sim 3$ ELMy H-mode cases (section 2.2.1). This means that the key to avoiding the onset of NTMs is not necessarily the exact β value [104], as one usually approaches the β-limit set by MHD, but it is the control of the size of seed islands and thus of intrinsic plasma perturbations. If these are smaller than the marginal island size, no NTMs are triggered (see figure 3). The main source of perturbations is the sawtooth activity, either the precursor/postcursor activity or the 1/1 mode responsible for the crash itself and its associated magnetic reconnection. It has been shown in ASDEX Upgrade [105] and DIII-D [106] that by avoiding sawteeth, with $q \geqslant 1$, the β onset could be raised by typically 20%, at which point NTMs are triggered by fishbones. On the other hand if the q profile is modified, as in $q > 1$ hybrid scenarios, because of the presence of impurities in the centre or for example due to ECCD current tailoring as shown in TCV, a seed island can be provided from a classical tearing mode with $\Delta' > 0$ [63].

However, a recent key result with respect to predictions for ITER shows that NTMs can be triggered at low beta, just above β_{marg}, when very long sawtooth periods are sustained using fast particles stabilization in JET [2, 37, 104]. As α-particles will strongly stabilize sawteeth (section 2.1.1), this is an important issue for ITER. However it has been shown that these stabilized sawteeth do not substantially deteriorate the overall plasma stability and control, but rather trigger 3/2 or even 2/1 NTMs at the sawtooth crash. An example from JET is shown in figure 7, where a relatively large island is triggered at the sawtooth crash directly with the width of the saturated island [4]. As these modes will grow in typically 30 s on ITER, they mainly cause concern because they can lead to the performance degradation.

The triggering of a 3/2 or 2/1 seed island by a sawtooth crash at $q = 1$ can result from global current profile modification [108] or from direct toroidal and non-linear coupling as obtained in 3D-MHD simulations [82]. The direct coupling of the 1/1 mode, causing the crash with 3/2 and 2/1 surfaces has been confirmed in MAST experiments where

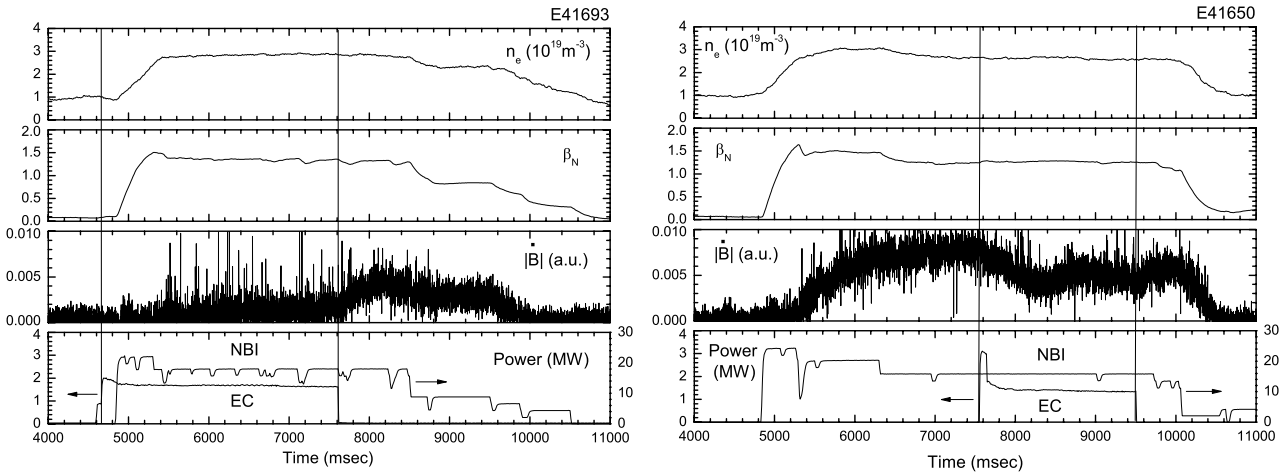

Figure 6. Application of early ECCD (lefthand figure) in JT-60U leads to a much smaller mode amplitude than application of the same ECCD power to a saturated mode (righthand figure) [102]. In pulse E41693 the EC injection angle is fixed so that the EC current is driven around the $q = 3/2$ surface. In pulse E41650 the injection angle is controlled in real time by searching the magnetic island position.

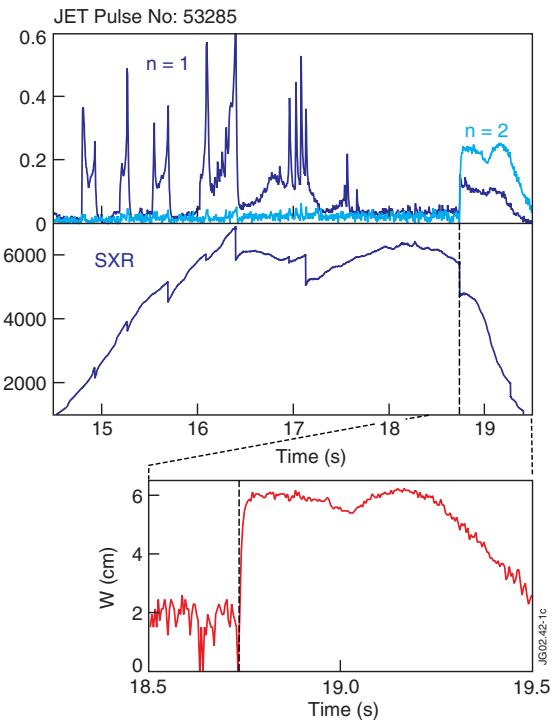

Figure 7. $m = 3/n = 2$ NTM triggered by a sawtooth crash being delayed (achieving a sawtooth period of 1.5 s) by the presence of ICRF fast particles at $I_p = 3.3$ MA, $B = 3.3$ T using minority H heating in the centre with $f = 51$ MHz. The island is directly triggered at a saturated size of 6 cm. Note that there is no significant $n = 1$ activity before the crash, and the relative SXR crash of about 15% is not much larger than the previous sawtooth crashes. (Reprinted figure with permission from [4]. Copyright 2002 by the American Physical Society.)

tight aspect ratio increases the coupling and NTMs have been triggered by sawteeth at low β [70]. This has also been used to explain the lower onset β observed in low q_{95} discharges [54], because of the larger $q = 1$ radius. Therefore one needs to destabilize sawteeth in order to avoid the large seed islands and this has been demonstrated on JET [4, 39], where significantly higher β onset values have been obtained using ICCD to destabilize sawteeth, albeit in scenarios without strong fast particles. The possibility of destabilizing sawteeth which are stabilized by fast particles, in scenarios similar to the ones expected in burning plasmas, has been demonstrated recently on JET [33]. This also confirms the possibility of decreasing δW_{fast} by increasing the shear at $q = 1$ (section 2.1.1).

Therefore in addition to direct control there is also the possibility of using ECCD on ITER to control the sawteeth and thus the seed islands triggering NTMs, as α-particle stabilized sawteeth are likely to be the main perturbation source. ECCD can be used to either stabilize or destabilize the sawteeth, by either slowing down or accelerating the build-up of the shear at $q = 1$ after a crash [14]. The best scenario would be to shorten the sawtooth periods in ITER, as sawteeth are helpful in avoiding impurity or He ash accumulation in the centre. However at present it is not known how to predict the level of destabilization required. Experiments are being performed to test this issue in ASDEX Upgrade [20], DIII-D and JET. Another possibility is to fully stabilize the sawteeth and thereby avoid NTMs onset, provided that fishbones or other perturbations are not sufficiently large to trigger NTMs, as is the case at high field and current in JET. However the relative size of the perturbations needed to trigger NTMs in ITER is expected to be much smaller than that in JET. Nevertheless, it is possible that the peaked electron heating provided by the self-heating mechanism will be sufficient to avoid impurity and He ash accumulation in sawtooth free scenarios in ITER.

Profile control and plasma shaping. According to the relevant contributions for the drive of NTMs as described in section 2.2.1, their growth and thus the resulting drop in confinement can be influenced by tailoring the radial profiles of current, pressure or density. The influence of the current and pressure profiles has been shown for example on JT-60U, where either the NTM's rational surface was shifted inside (resulting in both smaller local current gradients and pressure gradients [55]), or the pressure peaking was changed by using more central (negative NBI) or more off-axis peaked (positive

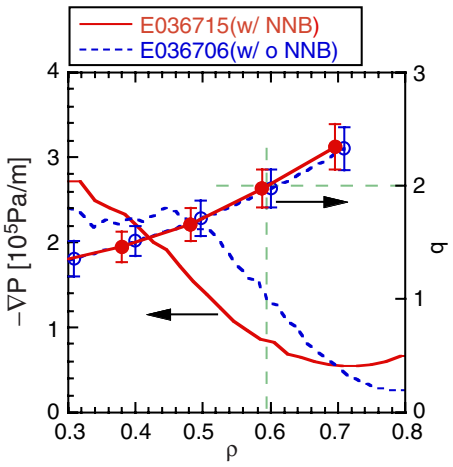

Figure 8. Two discharges from JT-60U with essentially identical safety factor (q) profiles measured by MSE. In the discharge with NNB injection a more centrally peaked pressure profile results and so the pressure gradient (∇P) is lower at $q = 2$ than in the discharge without NNB injection. This allows discharge E036715 to achieve a higher β_N without a 2/1 NTM being destabilized (which does occur in discharge E036706) [55].

NBI) heating [99]. Figure 8 shows an example from JT-60U [55] where negative ion neutral beam (NNB) injection is used to achieve a more peaked pressure profile with lower pressure gradient at $q = 2$. This allows a higher β_N to be realized in the an otherwise equivalent discharge with positive NBI heating (w/o NNB in figure 8). Plasma shaping can also have a direct effect via the tearing stability parameter Δ', as well as by its influence on the local pressure gradient: increased triangularity leads to higher pedestal pressure and thus to a smaller pressure gradient at the rational surface for a given β_N. In addition, the larger plasma current at given value of q_{95} causes a smaller bootstrap current fraction and thus a smaller NTM drive [109].

As the density gradient contributes more strongly to the bootstrap current than the temperature gradient, tailoring of the density profile is a very effective method to either avoid the NTM onset or to reduce its size. On ASDEX Upgrade density profile control has been demonstrated and is now used as a general tool to avoid NTM onset or to reduce its saturated size [38]. For weakly collisional plasmas ($\nu^*/\omega_D < 1$, ω_D: curvature drift frequency) density profiles tend to become peaked [110]. Central electron heating has been shown to be an effective tool to control this peaking as strong electron heating leads to an increased particle transport. This effect has been attributed to trapped electron modes being the dominant instability under such conditions [111, 112].

The frequently interrupted regime of NTMs (FIR-NTMs). As NTMs are driven by the loss of bootstrap current inside the magnetic islands, their width is expected to grow nearly proportional to the plasma pressure for given density and temperature profiles. On ASDEX Upgrade, however, a regime of (3,2) NTMs, the frequently interrupted regime (FIR), has been discovered [113, 114]. In this regime, the confinement degradation due to the (3,2) NTMs is strongly reduced in discharges with sufficiently high normalized plasma pressure at mode onset.

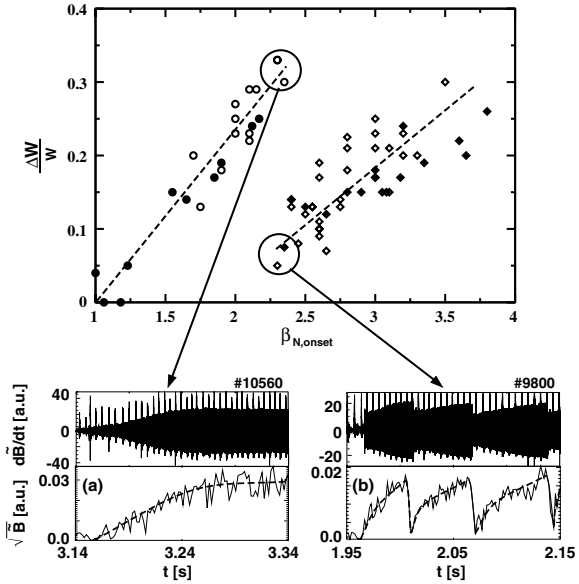

Figure 9. Comparison of reduction in energy confinement ($\Delta W/W$) due to (3,2) NTMs on ASDEX Upgrade (open symbols) and JET (full symbols). Very good agreement is seen, both in the relative confinement degradation as well as in the β_N value above which FIR-NTMs cause less energy losses. The lower figure shows the NTM behaviour for two ASDEX Upgrade discharges at about $\beta_N = 2.3$. The time-averaged amplitude for the FIR-NTM is significantly smaller (b) than the saturated amplitude of the smoothly growing mode (a).

In figure 9 the fractional confinement degradation of the plasma stored energy ($\Delta W_P/W_P$) due to the onset of NTMs is given as function of the β_N value at mode onset for ASDEX Upgrade and JET discharges, covering a wide range of plasma parameters (ASDEX Upgrade: $B_{tor} = 1.5$–2.5 T, $I_p = 0.8$–1.0 MA, $q_{95} = 3.0$–4.3; JET: $B_{tor} = 1.0$–2.7 T, $I_p = 1.0$–2.5 MA, $q_{95} = 2.5$–4.7) and for only dominantly NBI heated discharges. It is seen that for low values of $\beta_{N,onset}$ the confinement degradation increases linearly with the onset pressure as expected from NTM theory. At $\beta_N \approx 2.3$ the degradation in energy confinement is suddenly reduced (from about 30% to less than 10%). As the average H-factor at the mode onset for the discharges considered is often larger than 1, for β_N values between 2.3 and 2.45 the observed energy confinement is very close to that predicted by the ITER scaling.

As seen in figure 9(*a*), for low β_N values the NTM grows smoothly until it saturates at a finite amplitude. In discharges with $\beta_N, onset > 2.3$ however, the growth of the NTMs is often interrupted by sudden drops in NTM amplitude (see figure 9(*b*)). As the NTM growth time is larger than the time between two subsequent amplitude drops, the NTMs cannot reach their saturated amplitude. Thus, the time-averaged NTM amplitude is significantly reduced, reduced, resulting in a lower confinement reduction [114]. These sudden drops in NTM amplitude have been explained by non-linear coupling between the (3,2) NTM, (1,1) and (4,3) mode activity which are locked in phase [113]. The (4,3) mode typically only occurs at higher β_N and grows rapidly, suggesting it has ideal MHD origin.

It has been proven that by changing the linear stability properties of the ideal (4,3) mode, the critical β_N value for the transition to FIR-NTMs can be modified. Lower magnetic

shear at the $q = 4/3$ radius allows a transition into the FIR regime at smaller β_N values. This has been demonstrated on ASDEX Upgrade using localized co-ECCD just outside the $q = 4/3$ radius [115]. On JET a more global flattening of the central magnetic shear by early LHCD also resulted in strong FIR-NTMs [115].

Although the ITER operation point is planned to be somewhat below the observed $\beta_N \cong 2.3$ transition to the FIR NTM regime, it might be possible to consider operation at a little higher normalized plasma pressure. Thus, (3,2) NTMs might not be a very serious problem for ITER for $\beta_N > 2.3$; though current and density profiles may influence the exact β_N for transition to the FIR NTM regime. A weak magnetic shear that is characteristic of the recently developed advanced H-modes [106,116] would be additionally beneficial as it reduces the influence of the NTMs on confinement. The weak shear is destabilizing for the (4,3) mode and thus might decrease the critical β_N value for the FIR-NTMs. The reduced current gradient at the NTM rational surface has a stabilizing influence on the NTM.

Whereas in the discharges considered here, the non-linear coupling between a (3,2) NTM, (4,3) and (1,1) mode activity results in a reduction of the NTM amplitude, it has also been found recently that non-linear coupling of (3,2) NTMs to other MHD modes can have a similar effect. In [117] a reduction in the (3,2) NTM amplitude due to coupling to (4,3) and (7,5) modes with a corresponding improvement of confinement has been shown. Furthermore, large ELMs are also able to reduce the NTM amplitude [115]. The interaction between an NTM and external error fields of different helicity is able to affect small neoclassical islands [118]. This has been verified experimentally on DIII-D [93]. The non-resonant field control method for NTMs has the advantage that no precise positioning is needed and that it can also act on multiple possible NTMs simultaneously. The disadvantage found in the DIII-D experiments [93] was that the plasma rotation was greatly reduced (see section 2.4.1).

Conclusions. Besides active NTM stabilization by localized ECCD, as foreseen for ITER, in present-day experiments the onset of NTMs has been demonstrated to be influenced by seed island control, control of pressure and current gradients at the island's rational surface, as well as by plasma shaping. Seed island control allows for increasing the β_N value at NTM onset well above the marginal β_N value, whereas profile control and plasma shaping affect the NTM drive. For $\beta_N > 2.3$, a favourable NTM regime (FIR-NTMs) has been found, leading to only relatively benign confinement reduction due to (3,2) and (4,3) NTMs. If operated at sufficiently high normalized plasma pressures (as, e.g. anticipated for the hybrid scenarios), there might be no need to stabilize (3,2) NTMs in ITER. Thus, the effort for active stabilization might be predominantly concentrated on the more dangerous (2,1) NTM.

2.2.4. NTMs expected in ITER and their control.

Benchmarking NTMs for modelling of control in ITER. It should be possible to extrapolate from results of present devices, the saturated NTM island size and requirements on the ECCD power for NTM stabilization in ITER. The DIII-D discharge 114504 (figure 10), whose properties resemble

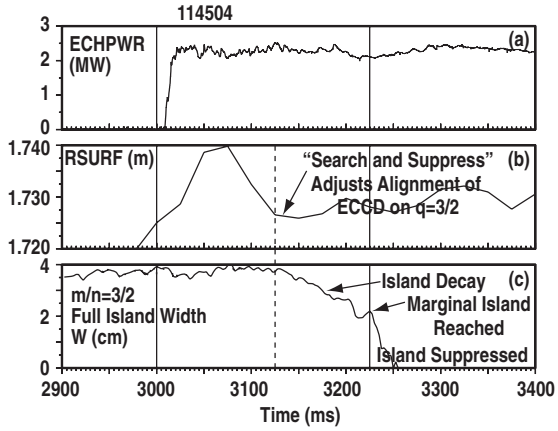

Figure 10. Detail of a DIII-D discharge, 114504, in which ECCD is used to suppress a 3/2 NTM. Showing (*a*) rf power, (*b*) plasma major radius adjusted to align $q = 3/2$ on ECCD, (*c*) $m/n = 3/2$ full island width calculated from $n = 2$ Mirnov signal.

ITER scenario 2 and which uses ECCD control to stabilize the 3/2 NTM, has been used to benchmark ECCD control in ITER Scenario 2.

Modelling is done by fitting to the modified Rutherford equation, equation (6) of section 2.2.1. For benchmarking to ITER, a form of the modified Rutherford equation is used where all current densities j_{bs}, $j_{||}$ and j_{ec} are flux surface averaged, and all scale lengths W, δ_{ec}, r_s and L_q are taken at the outboard midplane rational surface (here $j_{||}$ is the total parallel equilibrium current density, and j_{ec} is the peak rf current density). The modified Rutherford equation for NTM island growth or decay in this format is

$$\frac{\tau_R}{r_s}\frac{dW}{dt} = r_s\Delta' + a_2\frac{j_{bs}}{j_{||}}\frac{L_q}{W}\left[1 - \frac{W_{marg}^2}{3W^2} - K_1(W/\delta_{ec}, \Delta R/W)\frac{j_{ec}}{j_{bs}}\right] \quad (8)$$

with, for example in DIII-D, j_{bs} from ONETWO [119], $j_{||}$ from EFIT [120], j_{ec} from TORAY-GA [121] and a_2 expected to be a constant of order unity is fitted to the saturated no rf island width. For large islands the effective marginal island width is $W_{marg} \approx \sqrt{3}(W_{pol}^2 + W_d^2)^{1/2}$ which incorporates both the stabilizing effects of perpendicular thermal conductivity and polarization terms. $K_1 \propto \eta\delta_{ec}/W$ is the ECCD effectiveness numerically evaluated as a function of current drive width W/δ_{ec} and misalignment $\Delta R/W$ for unmodulated ECCD [122]. Parameters for the successful DIII–D ECCD stabilization of the $m/n = 3/2$ NTM are given in table 1. The inferred value before ECCD is $a_2 = 2.9$, using $r_s\Delta' = -m = -3$ and $L_q = 21$ cm. Given a_2 and W_{marg} then equation (8) determines the island evolution with ECCD with no free parameters. Note that the possible additional stabilizing effect of the ECCD reducing Δ' has not been directly experimentally verified to date and is not included here, though studies indicate this could be an important effect [123], as was demonstrated, in particular, in [90].

ITER Scenario 2, $m/n = 3/2$ NTM. ITER plans to have 20 MW of injected power at 170 GHz, from a launching point at $R = 6.485$ m, $Z = 4.110$ m. The relatively high

Table 1. Parameters for ECCD stabilization.

	DIII–D, $m/n = 3/2$	ITER, $m/n = 3/2$	ITER, $m/n = 2/1$
W_{marg} (cm)	2.5	3.2	4.3
δ_{ec} (FWHM) (cm)	2.7	10.0	8.2
W_{marg}/δ_{ec}	0.93	0.32	0.52
$j_{bs}/j_{\|}$	0.22	0.16	0.19
j_{ec}/j_{bc}	1.9	2.2	1.7

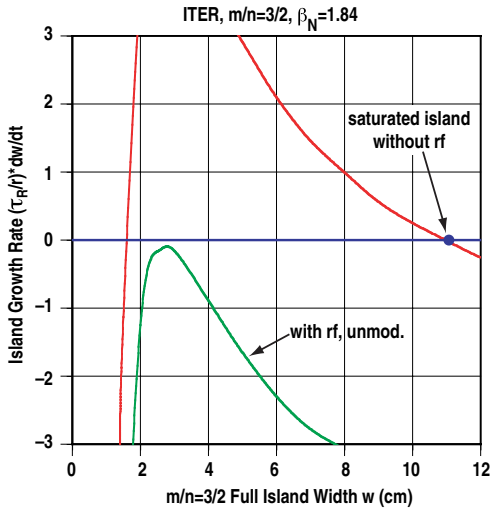

Figure 11. Island growth rate versus island width for the $m/n = 3/2$ NTM in ITER with no rf (red) and with just enough unmodulated rf power (39 MW) to achieve complete stabilization (green). Using 16 MW of modulated ECCD power instead would require operating at an island size where the green curve just touches $dW/dt = 0$. Parameters are $j_{bs}/j_{||} = 0.16$, $\Delta' r = -3$, $w_{marg}/\delta_{ec} = 0.32$, $\Delta R/w_{marg} = 0$, $j_{ec}/j_{bs} = 2.2$.

$(Z/a > 2)$ launching position makes the ray trajectories not very perpendicular to the $q = 3/2$ or 2/1 surfaces. Thus, narrow peaked current drive is difficult to achieve. Earlier work evaluated launching both at a somewhat lower elevation as well as from the midplane for a range of gyrotron frequencies [124]. In calculations, the tangential injection of EC beam was found most effective for the local current drive [125]. A simple criterion of peak ECCD current density comparable to the bootstrap current found that 20 MW of unmodulated injected power from an upper location from 170 to 210 GHz would be adequate at $q = 2$. Another study making different assumptions found that either 28 MW or 30 MW of EC power would be sufficient to suppress both $m/n = 3/2$ and 2/1 NTMs simultaneously [126]. This exceeds the planned 20 MW. Neither of these studies was benchmarked to actual successful experiments. However, experiments in JT-60U, extrapolated to ITER, concluded that 30 MW is required in ITER to stabilize both modes simultaneously without further optimization of the injection angles [127, 128].

The calculation of the modified Rutherford equation for the $m/n = 3/2$ mode for ITER from benchmarking to DIII–D experiments is shown in figure 11 for no ECCD and for j_{ec} (assuming $\Delta R/\delta_{ec} \equiv 0$) adjusted to give complete stability. For this case a_2 is taken from the DIII-D benchmarking and W_{marg} is deduced from the low $v/\varepsilon\omega_e*$ scaling of the polarization term [52] and by assuming convective parallel transport with Bohm-like perpendicular transport to give W_d [57]. Alternatively W_{marg} in ITER can be deduced from the observation that it varies as twice the ion banana width when data from DIII-D, ASDEX Upgrade and JT-60U are compared [123]; this scaling gives a smaller value for $W_{marg} \sim 1.8$ cm in ITER compared with value from the theoretical scaling, which is given in table 1 along with other parameters. Modulation of the ECCD as the O-point of the islands rotate by the launcher could increase the effectiveness about a factor of 2.5 [129]. This would reduce the required ECCD power but the island would have to be maintained near the peak of the green curve in figure 11 since modulation requires a non-zero island width for synchronization. However, this would represent a 75% reduction in island width which should substantially mitigate the effects of the NTM.

A narrower current drive in ITER would both increase j_{ec} and decrease δ_{ec} making the match to the 'marginal' threshold better and reducing the required power. A factor of two narrower ECCD deposition width would roughly increase the ECCD effectiveness a factor of two [122] and double the j_{ec} peak, which would result in about four times less ECCD power required for complete stabilization, i.e. reducing the power required from 36 to 9 MW.

ITER Scenario 2, $m/n = 2/1$ NTM. ECCD for stabilizing the $m/n = 2/1$ mode in ITER actually has two advantages over stabilizing the $m/n = 3/2$ NTM. First, the $q = 2$ flux surface is further out allowing the ECCD to point more perpendicular, thus making the current drive narrower. Second, the higher collisionality at $q = 2$, puts $q = 2$ part way into the high collisionality polarization threshold regime increasing W_{pol} and thus W_{marg}. Radial widths and current density ratios are given in table 1 and are taken from DIII–D experiments (a new value of a_2 appropriate to the m/n=2/1 is deduced by fitting DIII-D data) and from the ITER Scenario 2 equilibrium.

The calculation of the modified Rutherford equation for $m/n = 2/1$ in ITER is shown in figure 12 for no rf and for j_{ec} adjusted to give complete stability. Modulated ECCD power would reduce the required peak power but the island would have to be maintained at the peak of the green curve in figure 12. Again, improvements in the launcher to make narrower current drive can reduce the required ECCD power, i.e. a factor of 4 in power for a factor of 2 in current drive width.

Summary. The proposed 20 MW, 170 GHz ECCD system including the 'broad' current drive is adequate to substantially control the dangerous $m/n = 2/1$ neoclassical tearing mode and could be improved by 'narrower' current drive. Precise alignment on $q = 2$ and modulation of the peak ECCD power to drive current on the O-point and not the X-point is required to leave enough power (after mitigating the $m/n = 2/1$ instability) for $m/n = 3/2$ NTM stabilization and sawtooth control. Further common benchmarking of ECCD stabilization to experiments on ASDEX Upgrade, DIII-D, JET and JT-60U is in progress [123] and initial results support the conclusions given here.

2.3. Resistive wall modes

2.3.1. Physics of resistive wall modes. Steady-state high beta fusion plasmas with a high fraction of self-generated

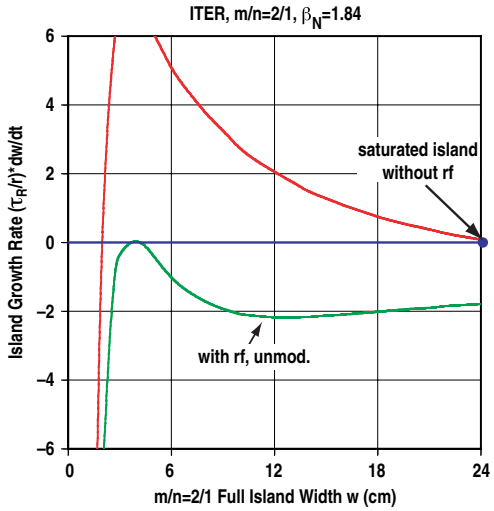

Figure 12. Island growth rate versus island width for the $m/n = 2/1$ NTM in ITER with no rf (red) and with just enough unmodulated rf power (25 MW) to achieve complete stabilization (green). Using 10 MW of modulated rf power instead would require operating at an island size where the green curve just touches $dW/dt = 0$. Parameters are $j_{bs}/j_{||} = 0.19$, $\Delta' r = -2$, $w_{marg}/\delta_{ec} = 0.52$, $\Delta R/w_{marg} = 0$, $j_{ec}/j_{bs} = 1.7$.

bootstrap current are likely to be unstable to long-wavelength external kink modes [1]. Such kink modes can in principle be stabilized by a nearby conducting wall, allowing beta to increase above the no-wall stability limit. However, if the wall has a finite conductivity, the growth rate remains positive and the kink mode is converted to a slowly growing resistive wall mode (RWM). By the time of the ITER Physics Basis [1], the possibility of stabilizing RWMs by plasma rotation in DIII-D discharges [130, 131], had been demonstrated. The experimental results had been fairly well reproduced by ideal MHD based models [132,133], which indicated that the critical rotation velocity, above which the RWM is stabilized, scales as a small fraction (a few percent) of the Alfvén speed. The computational predictions for the critical rotation tended to be somewhat higher than experimental observations, and the discrepancies were attributed to uncertainties in the drag, or dissipation, model [1]. The stabilizing drag on the RWM comes both from the Alfvén continuum damping and ion Landau damping. It had been pointed out that in a torus, trapping effects strongly reduce the ion Landau damping when the parallel phase velocity of the RWM in the plasma frame is subsonic [134, 135]. MHD computations with the MARS code, where the ion Landau damping was modelled as a parallel viscosity with an adjustable coefficient, showed that a rather large viscosity coefficient (about half the cylindrical value) gives critical rotation velocities close to experimental results [133].

Prior to the ITER Physics Basis [1], active feedback of the RWM had only been suggested theoretically. The main option at that time was a 'smart-shell' arrangement, with a large number of active coils covering most of the plasma surface. Subsequent developments in this area are discussed below.

Stabilization by rotation. Since 1999, the agreement between experiment and theory on rotational stabilization has been improved. A key element responsible for the improved agreement was the recognition of the important role of error fields [136, 137]. Garofalo *et al* [138 and 139], showed that the error fields slow down the plasma rotation, and also that reduction of the error fields by correction coils makes it possible to maintain the rotation and stability of discharges above the no-wall beta limit. The no-wall β-limit is important because advanced regimes in ITER are expected to operate somewhat above this limit, in the domain where RWMs can be unstable. Rotational stabilization of the RWM occurs up to the with-wall limit if the rotation is above a critical value which depends on β and q_{95}. At the with-wall limit an ideal instability co-rotating with the plasma occurs. The effects of β and rotation on RWM stability have all been observed experimentally in DIII-D and JET, where discharges were operated for a long duration over the no-wall limit, and produced measurements of the critical rotation speed [138, 140–142]. It has been pointed out that in steady state the error fields should be strongly amplified by plasmas near the no-wall stability limit [137] and [143] has treated the effect as being time dependent.

A computational study [144], using several different rotation profiles, indicated that for typical advanced tokamak profiles, the plasma rotation in the $q \approx 2$ region tends to be an important factor for stability. These calculations, where the damping was conservatively evaluated by using a small parallel viscosity, overestimated the critical rotation only by about 50% compared with experimental rotation profiles.

Further studies of the critical rotation velocity have been made using MARS [145, 146], modelling the ion Landau damping either as a parallel viscosity with an adjustable coefficient or by a semi-kinetic model without free parameters. In the semi-kinetic model, the drag is introduced as a local force on the displacement perpendicular to the field lines, derived from the imaginary part of the potential energy, δW_k, arising from kinetic calculations. δW_k was evaluated by an analytical large-aspect-ratio expansion [134] for circular plasma cross section, ignoring drift frequencies. An important conclusion from this model is that when the rotation is subsonic, the ion Landau damping comes predominantly from regions where the parallel phase velocity is near sonic, i.e. close to the resonant surfaces where $|m/q - n|v_{th,i}/R \sim \Omega_{rot}$ [145]. Here, toroidal trapping effects [134, 136] are insignificant. The parallel motion involves the toroidal side-bands; e.g. an $m = 2, n = 1$ magnetic perturbation couples to parallel flows with $m = 1, 3$, $n = 1$, and for small rotation speeds, these side-bands produce damping around the rational surfaces $q = 1$ and 3. Already when the rotation speed is a few percent of v_A, the kinetic damping is spread out across the entire cross section.

Modelling of the kinetic damping is work in progress. The fluid model with a small parallel viscosity usually overestimates the critical rotation, by the order of 50% [140]. The discrepancy can be removed by a larger parallel viscosity. However, there are experimental discharges that are manifestly stable [139, 145], where a large parallel viscosity fails to produce complete RWM stabilization. The kinetic model gives lower critical rotation that can be close to the experimental values, or underestimate them by as much as 40% [140, 142]. Kurita *et al* [147] showed that ferromagnetic walls give lower beta limits, but this effect will be small in ITER because of saturation.

Another important aspect of RWM physics, where experiment and theory can be compared, is the resonant field amplification (RFA), which is defined as the relative change in perturbed radial magnetic field from its vacuum value. Here, single-mode theory ignoring rotation [137] predicted that the steady-state RFA should peak (in fact, go to infinity [148]) at the marginal stability for a kink like perturbation, which was interpreted as 'the no-wall limit against the external kink' in [149], while in experiments with rotating plasmas the RFA increases significantly with pressure [141,150–152] above the no-wall limit. An alternative theory [143] shows that RFA is largest (though always finite) just at the stability boundary. Calculations with MARS using a fluid model with small parallel viscosity also makes the RFA peak, inconsistently with experiment, near the no-wall limit, while models with stronger damping (fluid model and cylindrical parallel viscosity or the semi-kinetic model) reproduce the experimental RFA for static error fields rather well [145, 151].

There are other aspects of the RFA experiments where the modelling is as yet less successful, e.g. the models in MARS tend to give too large real frequencies of stable RWMs. In [153], a single-mode model was used, where the plasma dynamics is represented by a single complex stability parameter, which determines both the RFA and the RWM growth-rate. According to the single-mode model, the experimental results in [153] imply that for stable RWM, the angular real frequency is typically τ_w^{-1} times a small number of the order 0.1–0.2, while MARS generally predicts larger frequencies. The semi-kinetic model currently implemented in MARS contains a number of approximations that can be removed; e.g. by including drift frequencies, kinetic modifications of plasma inertia, shaping effects on the damping, etc, and work in this direction is in progress.

Concerning scaling for extrapolation to ITER, both the fluid and the semi-kinetic model ignoring drift frequencies imply that the critical rotation velocity scales as the Alfvén speed (or equivalently, at fixed beta, the ion sound speed), and this agrees with DIII-D experiments [140]. (Among the approximations in the semi-kinetic model, the neglect of the drift frequencies may be the most important to remove, as it will give deviations from the Alfvénic scaling.) In DIII-D, the critical rotation speed at $q = 2$ is typically of order 1% of the Alfvén velocity (or less in recent experiments [154, 155]), and this is predicted to be marginally achieved by transport calculations for ITER with the ASTRA code [156]. Therefore, the strategy for ITER is to develop active feedback stabilization and also study its interaction with plasma rotation.

Active feedback. RWM feedback systems use magnetic sensor (and sometimes other) signals to gain information on the amplitude and phase of the unstable RWM, and apply this signal to control currents in a set of non-axisymmetric feedback coils. Models have been developed using ideal MHD. These models combine the force balance equation in the plasma ($-\nabla p_1 + \boldsymbol{j}_1 \times \boldsymbol{B}_0 + \boldsymbol{J}_0 \times \boldsymbol{b}_1 = 0$) for static equilibria, where inertia can be ignored) with $\boldsymbol{j}_1 = 0$ in the vacuum region and jump conditions for $\boldsymbol{b}_1 \times \boldsymbol{n}$ on the walls, which are modelled as thin resistive shells.

To determine the closed-loop growth rates, following practice in control theory, a non-dimensional transfer function $P(s)$ [157] is defined, giving the normalized flux or the magnetic field at the sensor position, resulting from a current in the active coils:

$$P(s) = b_s(s)/b_0 I_f(s). \tag{9}$$

Here, $b_s(s)$ is the field measured by the sensors, $I_f(s)$ is the current in the feedback coils and s is the Laplace transform variable. The normalization constant b_0 can be chosen as the field at the sensors produced by a 1 A current in the active coils in the absence of the walls and the plasma. In the simplest feedback system, the magnetic field detected by the sensors is fed back with a negative multiplier $-K(s)$, where $K(s)$ is the gain. The characteristic equation for the closed-loop is

$$1 + K(s)P(s) = 0. \tag{10}$$

In actual control systems using voltage control, it is useful to introduce another transfer function that measures the 'loaded' self-inductance of the feedback coil [157]. But for the control system currently envisaged for ITER, the active coils are far from the resistive wall, and voltage control gives only minor modifications over current control [158].

The simplest theories consider cylindrical geometry and assume that active and sensor coils can interact via a single toroidal and poloidal mode (m, n) [159, 160]. Although single-mode models ignore the geometry of the coil system, they give valuable insights into the main properties of the different types of sensors. Three main types of sensors have been analysed [145, 161]: radial, poloidal and radial sensors with the direct field from the active coils subtracted [162] (including the ac screening effect of the wall). One drawback of radial sensors located on the wall is that the signal from the unstable mode vanishes as the RWM approaches the ideal-wall stability limit [145, 163]. Poloidal detectors inside the first wall avoid this difficulty [145, 157, 161, 164, 165]. In the single-mode approximation, transfer functions for radial and poloidal sensors for a mode with poloidal mode number have the form [158]

$$\{P_m^r(s),\, P_m^p(s)\} = |m|(r_w/r_f)^{|m|-1}\{1,\, (2\Gamma_m + |m|)/m\}/(s - \Gamma_m) \tag{11}$$

with a single pole at $s = \Gamma_m = \gamma_m \tau_w$ which is the normalized growth rate of the RWM in the absence of feedback. Equation (11) shows that when the system approaches the ideal-wall stability limit, $\Gamma_m \to \infty$, the transfer function for radial sensors vanishes, while that for poloidal sensors remains finite.

Multimode cylindrical theories [145, 158, 166] take into account the shape of sensor and active coils by decomposing the coil currents into poloidal Fourier components. The total transfer functions are sums over single-m transfer functions:

$$P(s) = \Sigma_m c_m P_m(s) s_m, \tag{12}$$

where the c_m and s_m are geometrical coefficients characterizing the feedback and the sensor coils, respectively.

The multimode theories show that radial sensors fail when Γ_m is large [145, 158, 163], because of interference from the stable RWM, although ac compensation somewhat improves the performance of radial sensors [145, 161, 162]. Coupling to the stable RWM is not a major problem for poloidal sensors.

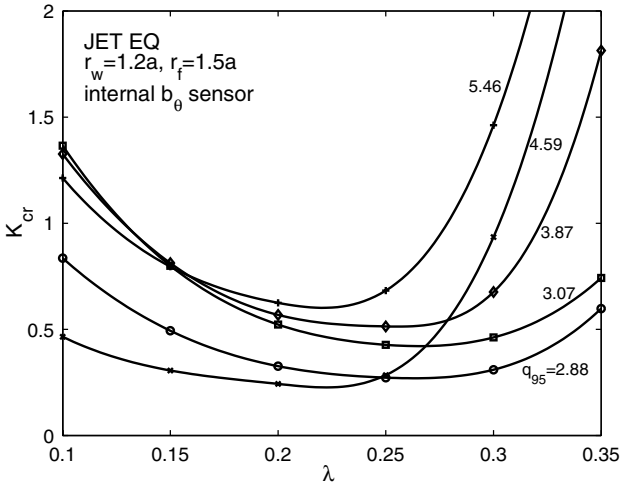

Figure 13. Critical gains as functions of the poloidal width of the coil for equilibria with pressure half-way between the no-wall and the ideal-wall limits and different plasma currents. The coil width is given by λ which is the fraction of the poloidal circumference subtended by the active coils (from [146] 'Robust control of resistive wall modes in tokamaks').

Even the double wall structure of ITER can be readily dealt with by a system using poloidal sensors [158, 167, 168].

The most realistic modelling of active feedback so far has been made with toroidal stability codes: MARS-F [145, 157, 168–170], the combination of DCON for the plasma and VACUUM for the vacuum-resistive wall region [164, 165], the KINX code [143, 171] and DCON combined with VALEN for three-dimensional modelling of the resistive wall [164]. Currently, benchmarks between these codes are in progress.

Also for toroidal systems, the transfer functions are sums over infinitely many isolated poles [165, 168], corresponding to the RWM growth rates in the absence of feedback. The transfer functions can be well approximated in the entire unstable half-plane by lumping several poles, and a three- (or for poloidal sensors even two-) pole approximation is generally very accurate [168].

Calculations with several codes [157, 164, 165] have shown superior performance of poloidal sensors. Work with MARS-F has shown that the control can be made robust with respect to variations in plasma pressure, current and rotation [145, 158, 168], if the active coils have a single coil in the poloidal direction located at the outboard midplane (e.g. as the C-coil in DIII-D). This is because such coils couple well to the high-beta $n = 1$ kinks which balloon on the outboard side of the torus. Figure 13 shows how the minimum proportional gain needed for stabilizing equilibria, about half-way between the no-wall and the ideal-wall limits, depends on the poloidal width of the coil for different plasma currents (q_{95}-values). A single coil array covering about 20% of the poloidal circumference [157] works well over a large range of plasma currents. Multiple active coils in the poloidal direction are useful when radial sensors are used [146, 164, 165], but do not bring substantial improvements in poloidal sensors [170].

The combined effect of feedback and rotation has also been studied [146, 168]. These studies show a synergistic effect if the toroidal phase of the feedback is shifted so that it pushes the mode in the same direction as the plasma rotation.

In summary, a simple coil system with poloidal sensors and a single feedback coil in the poloidal direction allow robust control with respect to changes in plasma pressure (up to a certain limit), plasma current and rotation, and can readily deal with the double wall structure for ITER [146].

2.3.2. Control of RWMs. As summarized in section 2.3.1, theory predicts that the RWM can be stabilized by rotation of the plasma of order of a per cent of the Alfvén velocity or by feedback control with non-axisymmetric coils or by a combination of rotation and feedback. At the time of writing of the ITER Physics Basis [1], early experiments in DIII-D [130] and PBX-M [172] had demonstrated that stable plasmas could be achieved with beta above the no-wall beta limit, by means of rapid plasma rotation driven by neutral beam injection. Such discharges were sustained at beta above the no-wall limit for up to 50 times the resistive magnetic field penetration time of the wall, demonstrating the principle of wall stabilization. Since then several other tokamaks, including NSTX [173], JT-60U [174] and JET [175], have also achieved beta values above the calculated no-wall stability limit. The critical rotation velocities for stability are consistent with the expected values of O(1)% of the Alfvén velocity [176]. However, in these experiments, extension to longer pulses proved difficult. Typically the plasma rotation became slower as beta exceeded the no-wall limit, eventually allowing the growth of a resistive wall mode.

Significant progress has been made in the last five years towards the goal of sustained wall stabilization of high-beta plasmas, using both rapid plasma rotation and direct feedback control. Recent experiments have developed techniques to sustain rotational stabilization for long pulses. Feedback control using non-axisymmetric coils has proven valuable in combination with plasma rotation, and initial experiments in feedback control with little or no rotation appear promising.

An important recent discovery related to RWM control has been the 'error field amplification' effect [136, 137, 177–179], also known as 'resonant field amplification.' This phenomenon is a consequence of the fact that plasma rotation provides only weak stabilization and rotation of the RWM, so that a small force can produce a relatively large (but stable) displacement. Simple lumped-parameter models for resistive wall mode stability [138, 150, 160, 162, 180] lead to an equation of the form

$$\tau_w \, dB_m/dt - \gamma_0 \tau_w B_m - H B_{ext} = 0, \qquad (13)$$

where B_m is the amplitude of the total perturbed magnetic field evaluated at the resistive wall, B_{ext} is the amplitude of an externally applied field that is resonant with the plasma mode, γ_0 is the growth rate of the mode in the absence of an externally applied field, τ_w is the inductive time constant for the resonant field to penetrate the resistive wall (in the cylindrical approximation $\tau_w = \mu_w \sigma \, dr/2m$, where σ, d and r are the conductivity, the thickness and the radius of the wall and m is the poloidal mode number) and $H = (1 + C\gamma_0\tau_w)/(1 - C)$. Here C is a factor that expresses the relative inductive coupling of the plasma (p), wall (w) and coils (c): $C = 1 - M_{pw}M_{wc}/L_w M_{pc}$, where M and L are the mutual and self-inductances. In the simple cases of a slab or a cylindrical

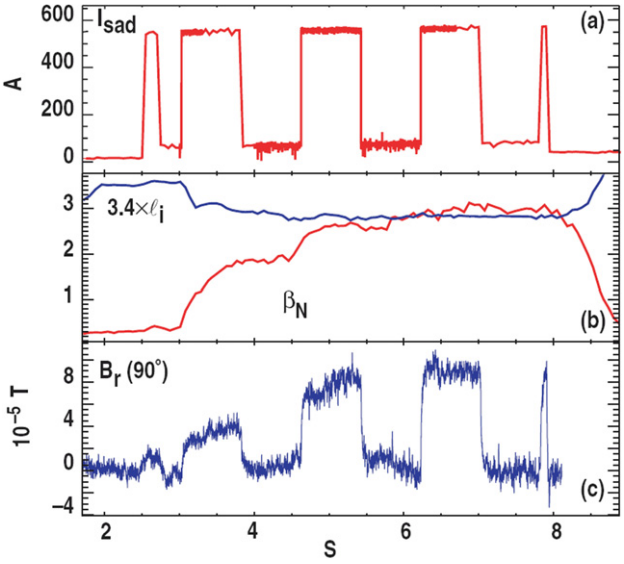

Figure 14. Discharge #59223, an error field amplification pulse in JET, showing time evolution of (*a*) current in the internal saddle coils, (*b*) normalized beta compared with $3.4l_i$, and (*c*) plasma response measured at the wall [175].

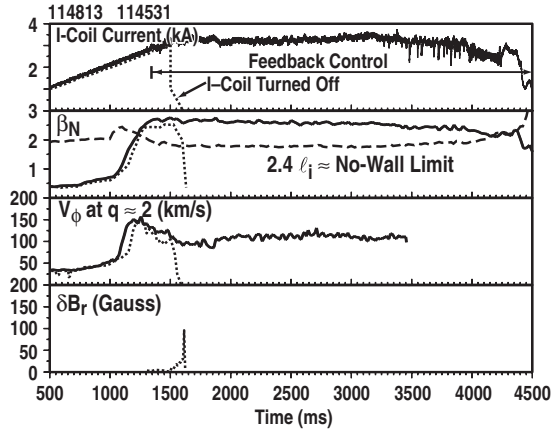

Figure 15. Feedback-controlled error field correction in DIII-D (discharge 114531, solid line) and a similar discharge without error correction (114813, dotted line). Shown are (*a*) I-coil current, which is feedback-controlled after $t = 1300$ ms in discharge 114531, (*b*) normalized beta and the estimated no-wall limit, (*c*) toroidal rotation speed at $r/a = 0.6$ and (*d*) amplitude of the RWM. (Reprinted with permission from [190]. Copyright 2004, American Institute of Physics.)

geometry with external coils, $C \to 0$ and H reduces to unity. In steady state, a weakly damped mode ($\gamma_0 < 0$) has a finite-amplitude response $B_m = (-\gamma_0 \tau_w)^{-1} H B_{ext}$ to a resonant external magnetic perturbation such as a magnetic field error. This enhanced plasma response, which becomes larger as the mode approaches marginal stability ($\gamma_0 \to 0$), can cause strong damping of the plasma rotation. The slowing of the rotation can then lead to loss of stability, as observed in DIII-D [181] and more recently in NSTX [173] and JET [142, 175]. (In contrast, a strong slowing of rotation is not observed at high beta in JT-60U [182], perhaps because the resonant component of the error field is smaller.)

Resonant field amplification has been demonstrated explicitly by the application of pulsed magnetic perturbations using non-axisymmetric coils [141, 175, 183–185]. Figure 14 shows an example from JET in which a series of square pulses was applied using internal saddle coils [175]. Here the no-wall ideal MHD stability limit is approximated by $\beta_N = 3.4 l_i$ (here the internal inductance of the plasma is $l_i = 2V\langle B_p^2\rangle/(\mu_0^2 I_p^2 R)$ with V = plasma volume), as confirmed by stability calculations and measurements [141]. The plasma response, measured at a location far from the driving coils, is almost undetectable at low beta, but increases strongly as beta approaches and exceeds the no-wall stability limit, where the RWM is only weakly damped.

The new understanding of resonant response to error fields has been used to improve performance in the wall-stabilized regime, and the feasibility of stabilization by plasma rotation now seems well established. Feedback-driven control of error fields has been effective in reducing magnetic field asymmetries in DIII-D, allowing continued rotation above the critical value for RWM stabilization. The resonant response of the stable RWM enhances the detection of small error fields, and the feedback system then corrects the error field by minimizing the enhanced plasma response [186]. This approach has allowed sustained operation above the no-wall limit using external control coils [183], and RWM stabilization by plasma rotation has become a routine element of DIII-D advanced tokamak experiments [187]. The recently installed internal control coils [188, 189] have provided similar performance [190], with wall stabilization sustained for more than 2.5 s (see figure 15). (In this discharge, a very broad current density profile and low triangularity lead to a relatively low beta limit; ideal MHD stability calculations show that the no-wall stability limit is well approximated by $\beta_N = 2.4 l_i$ with $l_i = \langle B_p^2\rangle/\langle B_p\rangle^2$). Similarly, reduction of the intrinsic field errors in NSTX by an adjustment of poloidal field coil positions has led to improved values of both toroidal beta and normalized beta [191, 192]. The low-activation properties of ferritic steel are attractive for fusion reactors, but its ferromagnetism introduces the possibility of additional field errors (see also section 2.4.1). Recent encouraging results from JFT-2M indicate that a ferritic steel wall can provide stabilization at high beta [193], although modelling predicts some reduction in the ideal-wall limit [194]. In both DIII-D and NSTX, a fast beta collapse or disruption with an ideal time scale magnetic precursor has been observed in some discharges, consistent with stability calculations indicating that in these cases beta is near the ideal MHD wall-stabilized beta limit [139, 192].

Despite the success of rotational stabilization, it is also necessary to develop direct feedback control of the RWM with non-axisymmetric coils. Burning fusion plasmas, lacking the large neutral beam torque of present experiments, may not reach the critical rotation frequency for RWM stabilization. Modelling with the MARS code [132, 133, 195] indicates that in the regime of subcritical plasma rotation, both feedback control and rotation can contribute to stability [169, 196]. As the rotation becomes small, stabilization must rely on feedback alone.

Recent experiments have yielded encouraging results for the feasibility of feedback stabilization of the RWM. Experiments in HBT-EP have shown that RWM instabilities related to current-driven kink modes can be stabilized with

direct feedback control [197, 198]. Feedback control of the RWM at high beta has been demonstrated in DIII-D [150, 162], using a set of six external control coils originally designed for error field correction. The performance of the feedback system is in good qualitative agreement with lumped-parameter models [138, 150, 162, 180]. Here equation (13) becomes

$$\tau_w \mathrm{d}B_m/\mathrm{d}t - \gamma_0 \tau_w B_m - G_m(\omega) H B_m = 0, \qquad (14)$$

where the feedback-controlled external field is $B_{ext} = G_m(\omega) B_m$. (In terms of the discussion in the preceding section, this equation is equivalent to the characteristic equation (10) combined with the single-mode transfer function (11). Here the time derivative is written explicitly.) The stability of the system depends strongly on the mutual inductive couplings of the sensors, mode, wall and external coils (the factor H defined above) and on the response of the feedback system, which is expressed here as the effective 'gain' function $G_m(\omega)$. (The function $G_m(\omega)$ includes the transfer function of the sensors, the gain of the control system and the response of the coil–amplifier system.) For example, at high frequency, induced currents in the wall reduce the coupling of the mode to radial field sensors but not to poloidal field sensors inside the wall, requiring the use of a larger gain with radial field sensors as γ_0 increases. The stronger coupling of poloidal field sensors may be particularly advantageous in multimode systems, where coupling to stable modes must be minimized. Indeed, experiments show that feedback control is improved by the use of internal poloidal field sensors [161] over radial field sensors, as predicted by numerical modelling [157, 158, 199] and analytic theory [161, 163, 180].

The internal control coils [188, 189] have further improved DIII-D feedback performance [190], providing stability against RWMs with higher growth rates. Internal coils are more effective because even at high frequency, their field is not completely cancelled by induced wall currents, allowing a faster response to rapidly growing RWMs. This can be expressed quantitatively using the coupling factor defined above: $C = 1 - M_{pw} M_{wc} / L_w M_{pc}$. In order for feedback control to provide the same stabilization as an ideal wall, the condition $C \geqslant 0$ must be satisfied [159, 200]. That is, the coupling of the coils to the plasma via induced wall currents must not exceed the direct coupling of the coils to the plasma. In a cylindrical model with helical currents, the case of external coils is just marginal: $C = 0$, while internal coils give $C = 1 - (a_c/a_w)^{2m} > 0$, where a_c and a_w are the minor radii of the coils and the wall and m is the poloidal mode number. The improved performance in DIII-D is qualitatively consistent with modelling using the VALEN code [164], which has predicted that in the actual experimental geometry, without the benefit of plasma rotation, internal control coils should allow feedback stabilization essentially to the ideal wall-stabilized limit. As shown in figure 16, internal coils have provided stability at higher beta and lower plasma rotation than was possible with external coils. In the figure, the trajectories of several experimental discharges in beta and rotation are compared with a representative stability boundary predicted by the MARS [132, 195] code in the absence of feedback [133, 169]. The MARS calculation uses parameters typical of DIII-D but not specific to these discharges, so the comparison is only qualitative. Nevertheless, the experimental results are reasonably consistent with the MARS prediction: a discharge without feedback becomes unstable near the calculated stability boundary, while discharges with feedback are able to cross the calculated boundary and continue to significantly higher beta and lower rotation.

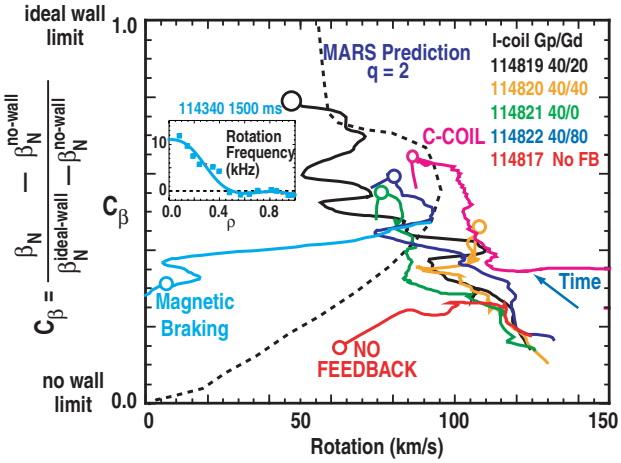

Figure 16. Trajectories of several DIII-D feedback-stabilized discharges in beta and rotation (at $r/a \sim 0.6$), compared with a representative stability limit predicted by MARS. Here beta is plotted as C_β, the fraction of the no-wall to the ideal-wall interval. A discharge without feedback is also shown. Open circles indicate the onset of the RWM. Inset shows the toroidal rotation velocity profile for the 'magnetic braking' case, shortly before the onset of the RWM. (Reprinted with permission from [190]. Copyright 2004, American Institute of Physics.)

The goal of future work is a full realization of the feedback control of resistive wall modes and confident extrapolation to burning plasmas. This will require validation of feedback models with and without plasma rotation [133, 169, 196, 199], using experimental plasmas with rotation rates above and below the critical value. Active MHD spectroscopy [183, 184, 201, 202] is a promising new technique for experimental investigation of the physics of rotational stabilization. Modelling indicates that a three-fold increase in the bandwidth of the DIII-D feedback system should allow feedback stabilization up to the ideal-wall beta limit in the absence of plasma rotation [190].

2.3.3. RWMs expected in ITER and their control. The RWM instability is expected in ITER steady-state scenarios with high values of β_N and low values of internal inductance. A typical representative of ITER steady-state scenarios is the Design Scenario 4 (see [203] where it is termed Scenario 1). This is a 9 MA highly shaped plasma with weak negative shear, producing about 300 MW of fusion power with $Q = 5$ for 3000 s ($R_p = 6.35$ m, $a_p = 1.85$ m, $\kappa_{sep} = 1.97$, $\delta_{sep} = 0.58$, $\beta_N = 2.57$, $l_i = 0.63$). The plasma cross-section is shown in figure 17. A set of Scenario 4-like plasmas was considered in the studies of RWMs expected in ITER steady-state scenarios. These plasmas have the same current, shape, about the same profile of q, but different values of β_N. To provide this β_N-scan, the plasma toroidal current was specified as

$$j_{tor} = j_0 \left\{ \frac{R}{R_p} \alpha G(\psi_p) + \frac{R_p}{R} [H(\psi_p) - \alpha G(\psi_p)] \right\}, \qquad (15)$$

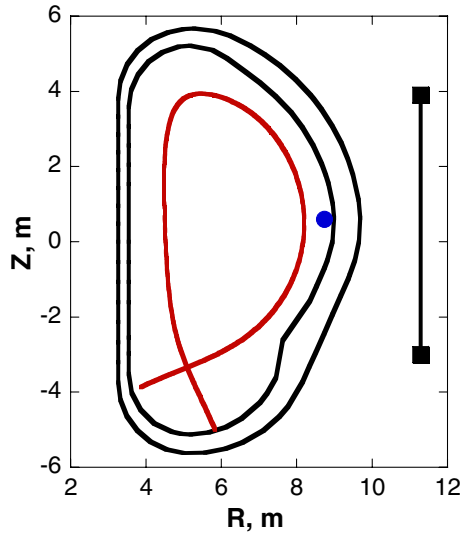

Figure 17. Plasma, vacuum vessel, magnetic sensor (circle) and coils(squares) used in the study of RWM control.

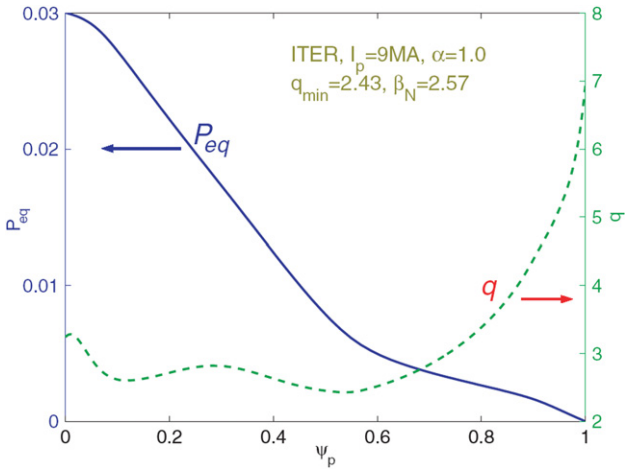

Figure 18. Profiles of plasma pressure, P_{eq}, and q for the plasma of Design Scenario 4 [146].

where j_0 is adjusted to get the plasma current of 9 MA, ψ_p is the normalized poloidal magnetic flux and the functions $H(\psi_p)$ and $G(\psi_p)$ [146] are obtained in the simulation of Scenario 4, corresponding to the case $\alpha = 1$, with the transport code ASTRA [203]. Different values of β_N are obtained by a variation in the parameter α. The profiles of plasma pressure, P_{eq}, and q are shown in figure 18. Table 2 shows the change in q and β_N with variation in α. For these the 'no-wall' β_N limit for $n = 1$ kink modes is about 2.4.

The main conducting structure affecting RWMs is the vacuum vessel consisting of the inner and the outer walls shown in figure 17. The vacuum vessel time constant, τ_w, for the mode $m = 1$, $n = 1$ is 0.25 s for an axisymmetric 2D model and 0.188 s for a 3D model, which takes into account the opening for the ports [204]. For Scenario 4 plasmas, the 'ideal-wall' β_N limit for $n = 1$ kink modes is about 3.6. The stability limits calculated for the 2D model of the vacuum vessel inner wall using different codes are given in table 3.

Table 2. Values of q and β_N as functions of the parameter α.

α	q_0	q_{min}	q_{95}	β_N
1.00	3.24	2.43	5.12	2.57
1.10	3.17	2.39	5.15	2.87
1.20	3.11	2.35	5.19	3.17
1.30	3.04	2.31	5.23	3.48

Table 3. The 'no-wall' and the 'ideal-wall' β_N limits for $n = 1$ kink modes.

Code	β_N (no wall)	β_N (ideal wall)
MARS-F [146]	2.45	3.65
KINX [207]	2.33	3.6
VALEN [208]	2.52	3.5
PEST-2 [209]	2.45	3.65

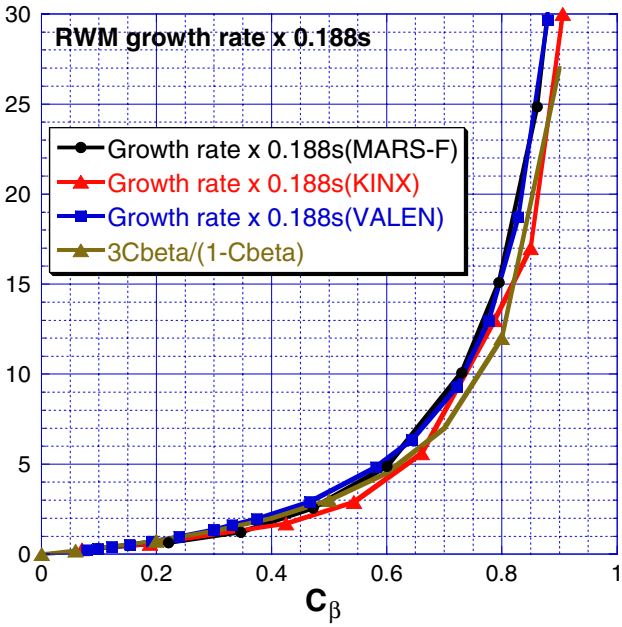

Figure 19. RWM growth rate normalized on $\tau_w = 0.188$ s versus C_β for ITER Scenario 4 plasmas calculated with codes MARS-F, KINX, VALEN and the analytical scaling. A 2D model of the vacuum vessel, no blanket.

The degree of RWM instability is characterized by a dimensionless parameter:

$$C_\beta \equiv \frac{\beta_N - \beta_N^{\text{no wall}}}{\beta_N^{\text{wall}} - \beta_N^{\text{no wall}}}, \qquad (16)$$

where $\beta_N^{\text{no-wall}}$ and $\beta_N^{\text{no-wall}}$ are the 'no-wall' and the 'ideal-wall' limits for the $n = 1$ kink mode. In Scenario 4, which has $\beta_N = 2.57$, $C_\beta \ll 1$ and only weakly unstable RWMs are expected.

Figure 19 shows the dependence of the RWM growth rate Γ on C_β calculated with the codes MARS-F (circles), KINX (triangles) and VALEN (squares) for the 2D model of the vacuum vessel inner wall. The solid smooth line shows the scaling $\Gamma\tau_w = 3C_\beta/(1 - C_\beta)$ [205]. The port openings reduce the 'ideal-wall' β_N limit by about 7% and increase the RWM growth rate by about a factor of 2, if the values of β_N between 3 and 3.4 are considered [204]. The RWM growth

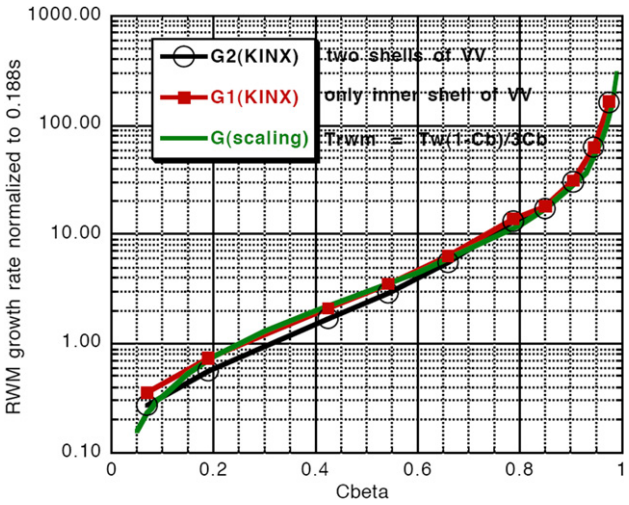

Figure 20. RWM growth rate normalized on $\tau_w = 0.188$ s versus C_β for ITER Scenario 4 plasmas calculated with code KINX for different models of the vacuum vessel: G2-two walls, G1-inner wall, G-scaling.

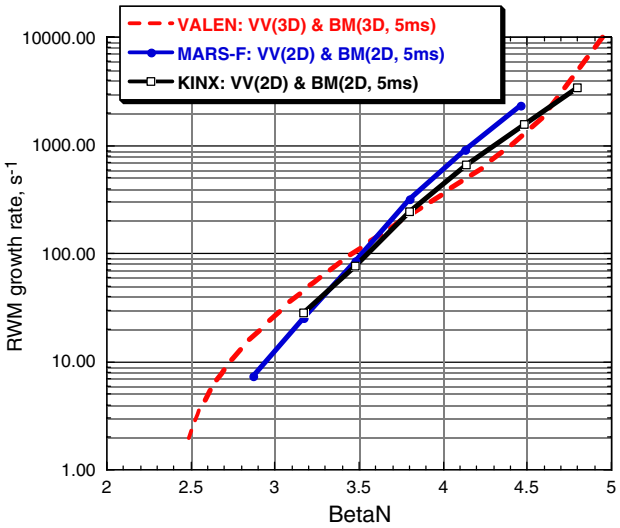

Figure 21. RWM growth rate versus β_N for ITER Scenario 4 plasmas calculated with codes MARS-F, KINX, VALEN. The blanket modules are taken into account. No plasma rotation.

rate as a function of C_β calculated with the KINX code with (the line with open circles) and without (the line with squares) account of the outer wall is shown in figure 20 [171]. The line without markers shows the scaling mentioned above. The outer wall only slightly reduces the RWM growth rate; however, it deteriorates RWM active stabilization using external feedback coils (see figure 24) by screening the magnetic field produced by the coils [167].

The blanket modules are located between the vacuum inner wall and the plasma. Each module is insulated from the other modules and consists of a thin first-wall panel and a thick shield block. The first-wall panel has many cuts, which strongly reduce the time constant of penetration of the magnetic field normal to the first wall. The time constant of a shield block is significantly higher, about 5 ms. The effect of blanket modules on the RWM growth rate was studied with the codes VALEN, MARS-F and KINX using very optimistic models, from the point of view of RWM stabilization. In the VALEN code, the blanket modules were represented by insulated conducting plates placed on the first-wall panel. Each plate has the time constant 5 ms. The codes MARS-F and KINX used axisymmetric models of the blanket modules. The blanket modules were represented by an axisymmetric conducting surface placed on the first-wall panels having the time constant 5 ms for the mode $m = 1$, $n = 1$. The dependence of the RWM growth rate on the value of β_N, calculated with these codes with the simplified models of blanket modules taken into account, is shown in figure 21.

The side correction coils (see section 2.4.2) may be used for the RWM control in ITER. There are three pairs of toroidally opposite coils connected to produce a magnetic field mainly with $n = 1$. Each pair has an independent power supply. The maximum coil current is 280 MAturns, the voltage insulation limit is 360 V/turns and the inductance of a pair of the coils is 50 μH/turns^2. About 120 MAturns are required for the correction of the error field expected in ITER (see section 2.5.2) and more than 160 MAturns are available for the feedback stabilization of RWM. At the position of the sensor shown in figure 17, a pair of coils produce the radial magnetic field of 0.1 T/MAturn. Each coil has 28 superconducting turns. The resistance of feedback circuit is determined by the resistance of busbars, 4 mΩ. This makes the time constant of the feedback circuit to be about 10 s, which is much higher than the vacuum vessel time constant τ_w. Therefore the resistive voltage drop in the busbars can be neglected in the studies of RWM control.

The magnetic sensors used in ITER in the feedback loop of the RWM suppression are located on the plasma side of the vacuum vessel inner wall. The position of the sensors, used in the studies of RWM control, is shown in figure 17. They measure the poloidal magnetic field of the modes $n = 1$. The expected amplitude of the noise in the sensors (signals other than RWM and ELMs) is less than 0.5 mT.

The study of RWM feedback control in ITER was done with the codes MARS-F and VALEN. The code MARS-F considers all RWM modes, but uses a 2D model of the feedback coils (effectively like a large number of coils) and a 2D model of the vacuum vessel. The VALEN code considers only a single mode (the most unstable), but uses more realistic 3D models of the feedback coils, the vacuum vessel and the blanket modules. Work to improve this modelling is in progress.

The studies of RWM control with MARS-F were done in three steps. Firstly, the transfer function, mapping the coil voltages to the magnetic field on sensors, was obtained. Then this transfer function was used for the design of controllers (second step). And thirdly, the controllers were validated in the simulations of stabilization of RWM. The stabilization starts when the magnetic field on the sensor is above $B_0 = 1.5$ mT. The studies have shown that, even with the double wall conducting structure, a simple coil system, with a single coil set in the poloidal direction, ensures robust RWM control with respect to changes in plasma pressure (up to a certain limit) and rotation. The results obtained for three controllers having different performances are presented in figure 22 [206]. The figure shows the maximum voltage required for the stabilization of RWM as a function of C_β. The voltage is proportional to B_0. The studies done with

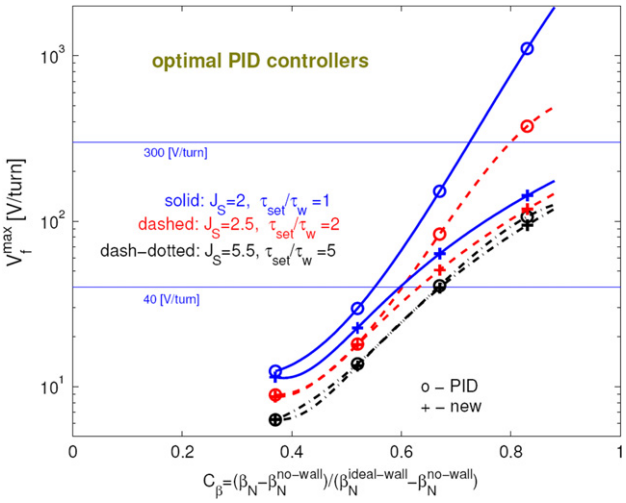

Figure 22. Maximum voltage required for stabilization of RWM by side correction coils versus C_β for three controllers with different performances. The stabilization starts when B_0 exceeds 1.5 mT. Code MARS-F and no plasma rotation [206].

MARS-F have shown that RWMs with C_β close to 1 can be stabilized, if the coil voltage is not limited. For stabilization of the highly unstable RWM, the feedback voltage should be proportional to a combination of the first and the second derivatives of the magnetic perturbation on the probes. Plasma rotation below the critical velocity does not significantly affect the performance of RWM feedback stabilization by the side correction coils.

The studies done with the code VALEN [204] show lower efficiency of the side correction coils with proportional and derivative feedback stabilization than was demonstrated in MARS-F simulations: RWMs with $C_\beta < 0.68$ can be stabilized. More detailed studies of the RWM stabilization are ongoing.

2.4. Error fields

2.4.1. Effect of error fields on plasmas.

Status as of the 1999 ITER Physics Basis. The effect of non-axisymmetric error fields in inducing locked modes in tokamaks was discussed in chapter 3, section 2.5 of the 1999 IPB [1]. The error fields occur due to poloidal and toroidal field coil imperfections and misalignments, from the current feeds to these coils and from ferritic material in the vicinity of the plasma. In the present experiments, to study their properties, error fields may also be deliberately applied using non-axisymmetric coil systems. Error field locked modes are operationally important because they tend to persist once established and then limit the performance or cause a disruption; conversely their correction can reduce disruptivity and allow lower density to be accessed [210, 211]. Error field locked modes are understood to arise from the braking torque applied to the plasma from a static error field, which can bring the rotating plasma to rest (most notably at $q = 2$) and allow islands to form (again most notably $m = 2, n = 1$).

In the 1999 IPB an empirical scaling was reported for the error field threshold $B_{\rm pen}$ (above which locked modes are induced) based on studies on JET, COMPASS-D [212] and DIII-D [213]:

$$B_{\rm pen}/B_{\rm t} \propto n^{\alpha_n} B_{\rm t}^{\alpha_B} q_{95}^{\alpha_q} R^{\alpha_R}. \qquad (17)$$

Data from Compass-D, DIII-D and JET were used to establish all the scaling coefficients except α_R which could then be determined from the Ohmic scale-invariance constraint, $8\alpha_n + 5\alpha_B - 4\alpha_R = 0$. There was agreement between the three machines that $\alpha_n \sim 1$, but $\alpha_B \sim -2.9$ on Compass-D and $\alpha_B \sim -1$ on JET and DIII-D. In addition in the IPB it was reported that the harmonic content of the error field (not just the magnitude of the $m = 2$, $n = 1$ component) affected the threshold, following from the previous studies [214, 215] and that as the ideal stability β-limit was approached, even in the presence of strong co-injected NBI heating (and therefore strong induced plasma rotation), the error field threshold decreased.

Since the IPB further error field experiments have been conducted on JET and DIII-D, and experiments have started on Alcator C-Mod [216] and TEXTOR [217] which amongst other things have confirmed the linear density dependence of the error field threshold. On Alcator C-Mod, which has linear dimensions of a factor of 4.3 smaller than JET, the typical locked error field threshold values, relative to toroidal field, are very similar to JET and DIII-D. Its threshold measurements also agree with the scaling based on scale-invariance arguments [1] to within a factor of 2, thus providing additional confirmation, based on direct size scaling. These further experiments on JET, DIII-D, Alcator C-Mod and TEXTOR have been aimed at refining the understanding of the physics of how error fields apply a braking torque to the plasma, additional studies of the effect of harmonic content and the effects of error fields near the $m = 2$, $n = 1$ NTM β-limit. In several cases these studies are still incomplete and so the results presented below represent a progress report.

Physics of error field braking. JET experiments [218] show a good match to theory, with locked mode formation being precipitated after the EM torque has slowed the plasma to approximately 1/2 its original frequency. Comparisons against various theory models indicate the ideal viscous model [219] best represents the JET data, with the error field threshold scaling with the unperturbed rotation as $\sqrt{\omega_0}$. In Alcator C-Mod clearly visible island structures occur in the electron temperature profiles at $q = 2$ when the locked mode is formed [216], confirming previous observations on JET [210]. The new JET results [218] have also shed light on the important physics of plasma rotation braking from applied error fields. A viscous drag model with the torque applied solely at the island ($q = 2$) surface would predict a uniform reduction of the rotation within $q = 2$, but the observations on JET indicate an approximately self-similar reduction within $q = 2$ (figure 23).

A new model associated with a toroidal viscous drag originating from non-axisymmetric fields, in particular due to the non-resonant $m = 0$, $n = 1$ mode, seems to qualitatively match these observations [220].

Experiments on DIII-D [93] have also shown an error field braking effect due to a largely non-resonant applied error field (dominantly $n = 3$, $m = 1$), and such an effect has also been discussed theoretically [222]. It is found that the $n = 3$ field

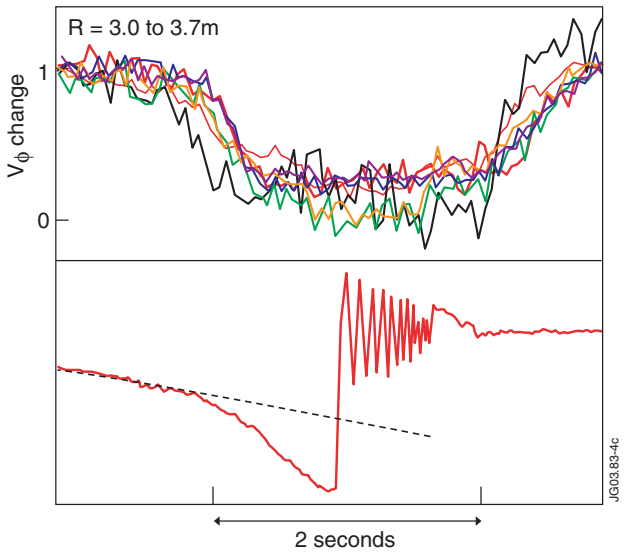

Figure 23. (Upper) Change in toroidal velocity for JET pulse 52067 (normalized to its value just before the error field is applied) for radii within $q = 2$. (Lower) Measured radial magnetic field. The error field locked mode is induced when this signal departs from a linear ramp (indicated by the broken line) expected from a direct response to the linearly ramped error field. As the error field is turned off, a slowly rotating island occurs, which slowly decays and rotates with increasing frequency. The plasma toroidal velocity then returns approximately to its initial value [221].

causes an 83% reduction of the momentum confinement time and that this is consistent with the expected braking from the $n = 3$ field ripple. However, in this case the $n = 3, m = 1$ (left- and right-handed) perturbation has $B_{r1,3}/B_t = 2.4 \times 10^{-3}$, compared with an otherwise equivalent case where an $n = 1$, $m = 2$ perturbation of only $B_{r2,1}/B_t = 4 \times 10^{-4}$ causes a 50% rotation reduction and subsequent locking [223] at similar β_N; this indicates that while these $n = 3$ non-resonant fields can indeed slow the plasma rotation, resonant fields have a greater drag at the same relative applied perturbation. Recent experiments have also shown the non-resonant effect on NSTX [224].

These results indicate that field ripple and neoclassical viscosity effects play a role in braking the plasma whereas previous models for locked mode thresholds, e.g. [225], have only considered the EM torque applied in the vicinity of the relevant resonant surface.

Effect of error field harmonic content. The poloidal mode coupling effect means that $n = 1$ harmonics, other than $m = 2$, affect the EM torque at $q = 2$. The EM torque applied at other resonant surfaces can also lead to a torque at $q = 2$ through perpendicular viscosity. Further, the neoclassical viscosity effects described above mean that non-resonant harmonics can lead to a torque.

At the time of the IPB in 1999, experiments on COMPASS-D [212] and DIII-D [215] had established empirical expressions for the torque (and thus the locking threshold) at $q = 2$ in terms of the $m = 1$–3, $n = 1$ harmonics. Subsequently the flexibility of varying the harmonic content using the six recently installed in-vessel coils on DIII-D, the I-coils, has been exploited to continue studies of error field harmonic effects [226, 227]. In these studies it is necessary to determine the intrinsic error field content (due to coil misalignments, etc) to add to the applied error field. Measurements in 2001 using an in-vessel rig determined the main sources of field error in DIII-D [228] and so the fields from these can be combined with those from the error field coils to give the total non-axisymmetric field. However, it should be noted that there are still difficulties in reconciling the applied fields required to empirically correct error fields with those determined using the measuring rig [227, 229].

The 1999 ITER Physics Basis expression for the effective error field from $m = 1$–3, $n = 1$ was [230]

$$B^2_{pen} = 0.28B^2_{r1,1} + (0.25B_{r11} + B_{r21} - 0.05B_{r31})^2 + 0.51B^2_{r31}, \quad (18)$$

where $B_{rm,n}$ is the unweighted (m, n) Fourier component evaluated in straight field line coordinates at $q = 2$. This form was based on the theoretically expected form, whereas a previous simpler empirical expression (equation (19)) has been used extensively in the design of the ITER error correction coils (see section 2.4.2):

$$B^2_{pen} = B^2_{2,1} + 0.8B^2_{r3,1} + 0.2B^2_{r1,1}. \quad (19)$$

Recent fits based on DIII-D I-coil data are given in [227]. Comparison of all the fits for 3-mode coupling against the DIII-D and the COMPASS-D data shows at present that there is no satisfactory fit covering both machines, possibly reflecting the different toroidal field scaling found in the two machines. Overall the simplest empirical fit (equation (19)) best describes the data in both machines. This '3-mode' criterion is applied to the design of the ITER correction coils (see section 2.4.2) where the requirement is that $B_{3\text{-mode}}/B_t < 5 \times 10^{-5}$ for avoiding locked modes. The choice of the constant (5×10^{-5}) is conservative relative to the predicted error field sensitivity of ITER [210].

Error fields due to ferritic material. In order to reduce ripple losses, regions of ferritic steel have been installed in stages in JFT-2M. First, ferritic plates (FPs) were installed outside the vacuum vessel. Ferromagnetism of the FPs could be a problem because of possible low-n error fields which might be induced. In order to examine the effects of error fields due to the FPs, experiments were performed with an unbalanced placement of FPs, where 40% of the FPs contiguous in the toroidal direction were removed, to increase the error field of the $n = 1$ mode artificially [231]. The error field with the unbalanced FPs was evaluated, and, for example, the magnetic field from FPs magnetized by the vertical field was estimated to be 1.5 mT. The evaluated error field was $B_{r2,1}/B_t = (0.6–1) \times 10^{-4}$ for B_t of 0.8–1.3 T. This value corresponds to the 20–40% of the critical value for the locked mode and so was not an operational issue for JFT-2M.

In the second stage, the inside wall of the vacuum vessel was fully covered with the ferritic steel (termed a 'ferritic inside wall' (FIW)) as a simulation of the blanket wall of a reactor, in JFT-2M. To study possible error field effects the entire FIW is assumed to be rigidly shifted by 20 mm, much more than the estimated actual shift of (~3 mm). The induced error field due to the shift of 20 mm, using the 3-mode coupling formula (equation (19)) is less than $B_{pen}/B_t \sim 4 \times 10^{-4}$. This is lower than the allowable limit for an Ohmically heated plasma

$B_{pen}/B_t = 5 \times 10^{-4}$, and experimentally the region of locked mode occurrence was not affected by the FIW. These tests thus increase confidence that there are no unexpected error field effects from ferritic materials, such as will be used in the test blanket modules in ITER (see section 2.4.2).

Other issues. Initially as β is increased with NBI heating the error threshold rises; this effect is studied in [218] where it is shown at low-β that the error field threshold rise with the beam induced rotation is best fitted as $\sqrt{\omega_0}$. In contrast, equivalent plasmas where solely ICRF heating was used in JET to increase β showed only a small increase in the error field threshold [210]; though it should be noted that recent results from TEXTOR do show an increase in the error field threshold with ICRF heating power [217]. At higher β as the ideal stability limit is approached the error field threshold decreases with rising β, as reported in the IPB. It is now understood that as the ideal β-limit is approached in low collisionality ELMy H-mode plasmas, $m = 2$, $n = 1$ NTMs are likely to be destabilized just below the ideal limit [232]. Studies are in progress on how the 2/1 NTM threshold is affected by error fields or conversely how the error field threshold changes as the 2/1 NTM limit is approached. Initial studies on DIII-D [58] indicate that error fields can substantially lower the 2/1 NTM β-limit (by $\sim$ 40%), with error fields causing rotating 2/1 NTMs to be formed (since they rotate this eliminates the possibility that they are error field driven modes). Initial analysis of similar studies on JET [58] indicates a reduction in the error field threshold near the 2/1 NTM β-limit, but there is no evidence of rotating modes being formed in this case. Previous JET experiments [210] have shown that error-field-induced modes can act as the seeds for 2/1 NTMs.

Studies on Alcator C-Mod [216] have extended the range of toroidal field over which the error-field-induced locked modes have been studied to ~8 T (i.e. encompassing the ITER range). This is important because the toroidal field scaling of the error field locked mode threshold is one of the remaining uncertainties. An identity experiment (same ρ^*, ν^*) with JET also shows good agreement on locked mode thresholds and further studies are ongoing in this area [216].

2.4.2. *Error fields expected in ITER and their correction.*

Expected error fields. Analysis of ITER error fields and their correction is reported in [233, 234]. The plasma corresponds to the initial phase of current flattop in the 15 MA inductive scenario (Scenario 2).

Misalignments of current centrelines of the toroidal field (TF) coils, central solenoid (CS) and poloidal field (PF) coils, arising during the coils manufacture, installation and assembly, are the main contributors to the error field. A tolerance analysis was used to estimate the expected deviation of the coil centrelines from the nominal position and shape. The deviations are in the range 2–3 mm. The Monte Carlo simulation, used for error field analysis, included 228 degrees of freedom in deviation of the current centrelines. The total number of simulations was 250 000, which guarantees reasonable accuracy of the results. The expected '3-mode' error field (equation (19)) due to misalignment of the PF, CS and TF coils is less than 11.9×10^{-5} with a probability of 99.9%. The contribution to error fields from the CS and PF coils is approximately the same as it is from the TF coils.

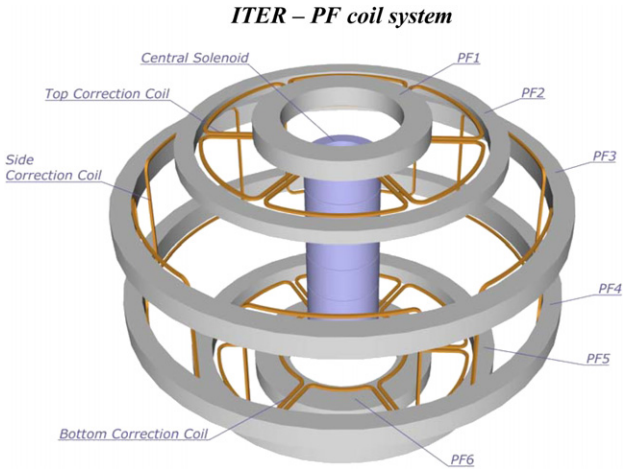

Figure 24. ITER CS and PF system with error field correction coils.

Error fields due to the CS, PF and TF coil joints, feeders and terminal regions were calculated using simplified models of the current-carrying parts of these elements. Each pair of TF coils gives an error field comparable to the error field from a PF coil. However, the symmetry of the TF coil systems provides a mutual compensation of the individual error fields generated by the coil pairs, so the total error field is negligibly low. The total '3-mode' error field from the elements of CS and PF coils is 3.0×10^{-5} of the toroidal magnetic field. The main contributors are the terminals, joints and feeders of six CS modules.

Six test blanket modules (TBM) are planned to be used in ITER. Five of them will have ferromagnetic structure material (martensitic steel). The '3-mode' error field from the five TBMs is 3.1×10^{-5}. However, when this error field is superimposed on the error field from the coil joints, feeders and terminal regions, the total '3-mode' error field is only 4.0×10^{-5}.

There are two neutral beam (NB) injectors and one diagnostic NB injector in the present design of ITER. To avoid deflection of the ion beam, the magnetic field inside the NB injector should be very low during its operation. Each injector has a magnetic field reduction system (MFRS) designed to reduce the tokamak poloidal field inside the injector to an acceptable level. The MFRS comprises active coils and passive ferromagnetic shields. The '3-mode' error field due to the NB MFRSs (two NB injectors and the diagnostic injector) in Scenario 2 is about 1.1×10^{-5}.

Error field correction The correction coils (CC) system has been designed for ITER to reduce the error fields to an acceptable level. The coils are shown in figure 24. The CC system consists of three sets of superconducting coils (top, side and bottom) distributed poloidally to provide control of several m mode numbers. Each coil is connected to the toroidally opposite coil in series, to produce an $n = 1$ asymmetric mode and to avoid coupling with axi-symmetrical magnetic fields. Three power supplies are used for each CC set to provide toroidal rotation of the correcting field.

Table 4. Currents (kAt) required for correction of ITER error fields and the nominal currents.

Source of error field	Top CC	Side CC	Bottom CC	Comment
NBI MFRSs, 5TBMs and Coil joints, feeders, terminals	20.6	12.5	20.0	Correction to 0
TF, CS and PF Coils misalignment	103.6	104.2	160.1	Correction from 11.9×10^{-5} to 5×10^{-5}
All sources	124.2	116.7	180.1	Simple summation
Nominal currents	140	280	180	

The currents required for correction of the error field expected at SOF are summarized in table 4. The table also shows the nominal currents in the CCs. The side CC have additionally about 160 kAt available for control of resistive wall modes.

2.5. MHD stability in advanced scenarios

2.5.1. MHD stability in plasmas with weak magnetic shear. At the time of the ITER Physics Basis [1], two types of advanced scenario had been identified, associated with a centrally peaked current density profile (high l_i) or a very broad current density profile (reverse shear). The high-l_i scenario had good confinement and a high-beta limit ($\beta_N \sim 4l_i$, with $l_i \sim$ 1.0–1.5), but because of high values of edge safety factor and poor alignment of the bootstrap current with the desired current density profile, it was considered a less than optimum candidate for steady-state operation. Reverse shear plasmas generally have better bootstrap current alignment for steady state, but somewhat lower l_i (∼0.7–0.8) and correspondingly lower β-limits without wall stabilization. Furthermore, discharges with strongly reversed central magnetic shear were often observed to form a strong internal transport barrier, leading to strong local pressure gradients which can further reduce the beta limit to $\beta_N \sim 2$. The limiting MHD mechanism in this case was an ideal instability, the infernal mode, with its mode structure having a maximum near the region of minimum safety factor (q_{min}). (In addition to these β-limiting ideal instabilities, in some cases, resistive neoclassical tearing modes were observed in regions of positive shear. TAE instabilities have also been observed.) By avoiding a strong internal transport barrier, a high-β scenario with weak central shear, broad pressure profile and wall stabilization of the $n = 1$ kink mode was identified experimentally. In this scenario, $\beta_N \sim 4.5$ was roughly 1.3 times the free boundary ideal MHD limit. This case was considered a promising steady-state candidate, but at that time it was uncertain whether wall stabilization could be sustained in long pulses (see section 2.3.2).

Recent experiments and modelling have extended the performance and the understanding of stability limits in the 'high-l_i' and 'reverse shear' regimes, and plans for ITER now include advanced scenarios related to these two regimes. In this section progress in the macroscopic stability of advanced regimes with weak central magnetic shear is discussed. The case with strongly negative shear is addressed in section 2.5.2. (Here central magnetic shear refers to the shear between the magnetic axis and q_{min}; weak shear scenarios typically have q_0 (safety factor on-axis) in the range $1 < q_0 < 2$, with q_{min} less than or equal to q_0.)

Discharges with centrally peaked current density profile are of interest because of their potential for high confinement and high-beta limits without the need for wall stabilization, but discharges with high l_i have drawbacks for steady-state operation. However, theoretical optimization studies [235] suggest the existence of a regime of moderate l_i, with q_0 near unity and low central magnetic shear, more suitable for steady state while retaining some of the advantages of a peaked current density profile. With a sufficiently broad pressure profile (pressure peaking factor $p(0)/ < p > \sim 3$) the beta limit in such cases can be raised to $\beta_N = 3.5$–4.0 without wall stabilization, through strong plasma shaping with triangularity up to 0.7. An additional feature of such a configuration is the relatively high-bootstrap current, with f_{bs} up to 70% and good alignment to the total plasma current density.

Candidates have been developed recently for high-fluence or 'hybrid' operation in ITER, building on earlier experiments in stationary H-mode discharges with improved core confinement [236–238]. These operating regimes have some features in common with the moderate-l_i regime described above. (Hybrid mode operation is discussed in more detail in chapter 6 of this issue [2].) Typically, a stationary q profile with low central shear is maintained with $q(0)$ and q_{min} just above unity, but with very small sawtooth oscillations or no sawteeth at all [99, 116, 239, 240]. The effective beta limit is set by the 2/1 tearing mode. However, in the absence of sawteeth this mode is encountered only when beta approaches the no-wall ideal kink mode limit at $\beta_N \geqslant 3$, a significantly larger beta than the typical threshold for 2/1 tearing modes in the presence of sawteeth. Avoidance of peaking of the current density profile leading to sawteeth is related to the presence of a benign 3/2 tearing mode [99, 240, 241] or in some cases a 1/1 fishbone instability [116, 242]. This has allowed, for example a quasi-stationary equilibrium with $q_{min} > 1$ to be maintained at $\beta_N \sim 3$ for over 6 s in DIII-D (see [241] and chapter 6 of this issue [2]).

The strong internal transport barriers associated with strongly negative magnetic shear can also occur in discharges with weak central shear. In such a case, the main stability limitation is set by the $n = 1$ global kink, driven by the combination of a peaked pressure profile and a weak shear. Typical beta limits of $\beta_N \sim 2$, set by a fast-growing $n = 1$ mode, have been observed in many experiments [243–246]. Analysis of experiments on JT-60U [244, 247] confirm that when q_{min} has a rational value, a variety of MHD events are triggered, including partial beta collapse, disruptions and energetic particle driven instabilities. Partial collapse is associated with localized spikes in the pressure gradient, and disruptions occur when there is a strong coupling to the edge, caused either by a low shear or a broader pressure profile [76].

Experiments have shown clearly that the beta limit due to the $n = 1$ ideal kink mode can be significantly improved with a broader pressure profile [243, 245] in good agreement with ideal MHD stability modelling predictions [245, 248]. The transition from an L-mode to an H-mode edge is an effective

means of broadening the pressure profile, but edge-localized modes (ELMs) that occur during H-mode often degrade the internal transport barrier [249]. Some experiments have reported a role for energetic particle driven modes, which may either trigger formation of an internal transport barrier [246] or trigger a transition to the H-mode [245]. With a broad H-mode edge pressure profile and wall stabilization, the $n = 1$ kink stability limit increases to $\beta_N \sim 4$, and in some cases significantly greater.

A promising related scenario is the quiescent double barrier (QDB) mode, achieved with counter neutral beam injection [250–252]. These discharges are in the H-mode but without the deleterious ELMs, allowing an internal transport barrier to be maintained. Regulation of the heating power allows the volume within the internal transport barrier to be expanded slowly, without creating local pressure gradients strong enough to destabilize the $n = 1$ kink mode [250]. In these discharges, normalized beta up to $\beta_N \sim 2.9$ is achieved. Full characterization of the global stability limits of the QDB mode remains an open issue.

Stability modelling and experiments since the ITER Physics Basis [1] have influenced the redesign of ITER with higher triangularity. In general, experimental results are in good agreement with predictions that the $n = 1$ ideal kink stability limit depends strongly on the discharge shaping, the current density profile, the pressure profile and the presence of a conducting wall, often in a synergistic way [248]. Numerous modelling efforts have shown that the β_N limit improves with triangularity, but only when given sufficient elongation [248, 253] and a sufficiently broad pressure profile [248]. Experimental results [254,255] are consistent with these predictions. As discussed above, the β_N limit varies inversely with the peaking of the pressure profile in both modelling and experiment [191, 243, 245]. Very high values of β_N have been achieved in recent spherical tokamak experiments [191, 256], consistent with modelling predictions on the β_N variation with aspect ratio [253].

The weak shear or reverse shear regime with low l_i and $q_{min} > 1$ which is proposed for steady-state advanced scenarios in ITER is likely to require wall stabilization. In such discharges the experimentally observed no-wall β_N limit decreases with increasing q_{min} [255], consistent with modelling predictions (see figure 25). Experiments are also consistent with the prediction that the beta limit can be significantly increased by wall stabilization, particularly in the case of a broad current density profile (high q_{min}) and a broad pressure profile. Although the cases shown in figure 25 with beta well above the no-wall beta limit are transient, other experiments have demonstrated the feasibility of long pulses with wall stabilization (See section 2.3.2).

2.5.2. MHD stability in plasmas with strong negative magnetic shear. Reversed shear discharges are characterized by negative magnetic shear in the core region, with the internal transport barrier near the minimum q region. Therefore, a large pressure gradient exists in a weak or negative magnetic shear region, and a large current density can exist near the plasma edge. These characteristics often cause harmful MHD instabilities over a wide range of β_N. It has been observed in reversed shear discharges that the resulting ITBs can lead

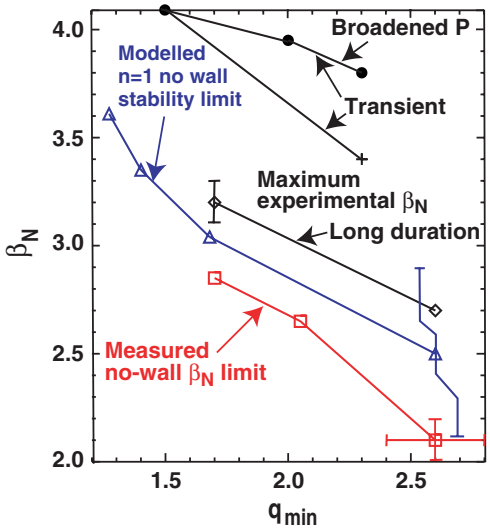

Figure 25. The measured and modelled dependence of the β_N limit on the minimum safety factor in DIII-D. Squares are the experimentally measured no-wall limit and triangles are the $n = 1$ no-wall limit calculated for equilibria with profiles similar, but not identical, to the profiles in the experimental discharges. Diamonds are the maximum β_N at which discharges have been operated for the duration of the machine pulse without significant instability, circles are the maximum β_N that have been obtained in steady-state scenario discharges. The cross represents the maximum β_N obtained without adding extra gas to broaden the pressure profile. (Reprinted with permission from [255]. Copyright 2005, American Institute of Physics.)

to strongly peaked pressure profiles and drive ideal kink instabilities which result in a disruption in JT-60U [174, 257], JET [245, 258], DIII-D [259] and ASDEX Upgrade [260]. Unique to the reversed shear regime is a dominantly $n = 1$ mode, which has multiple harmonics. This mode is a seemingly common limit to achieving the highest performance plasmas. Ideal MHD stability analyses revealed that the upper limit of the achievable β_N is in agreement with the calculated ideal stability limit of the $n = 1$ kink-ballooning mode [245, 257–259, 274]. The region of major collapses, or disruptions, in reversed shear discharges in JT-60U is shown on the $\beta_N - q*$ plane in figure 26. Here, q^* is the safety factor near the edge defined as $q^* = (5a^2B_T/2RI_p)[1 + \kappa^2(1 + 2\delta^2 - 1.2\delta^3)]$, here κ and δ are the elongation and the triangularity of the plasma poloidal cross-section, respectively. The phenomenology of these disruptions is identical to those found in low shear advanced regimes.

In DIII-D, an L-mode negative central shear (NCS) discharge exhibited a strong ITB and continually increasing pressure peaking. Early calculations suggested the MHD burst prior to the disruption correlated with a resistive interchange instability; this instability was believed to reduce core rotation thereby triggering a more global double tearing mode that was thought responsible for the actual disruption. Re-analysis of the discharge just before disruption, however, found that a small (∼15%) steepening of the pressure gradient inside $\rho \sim 0.3$ can lower the predicted ideal beta limit to that of the observed disruption limit, implying that the disruption can be consistent in this case with an ideal instability [259]. A new quasilinear ideal MHD stability model [261], based on the evolution of a discharge through an ideal stability boundary,

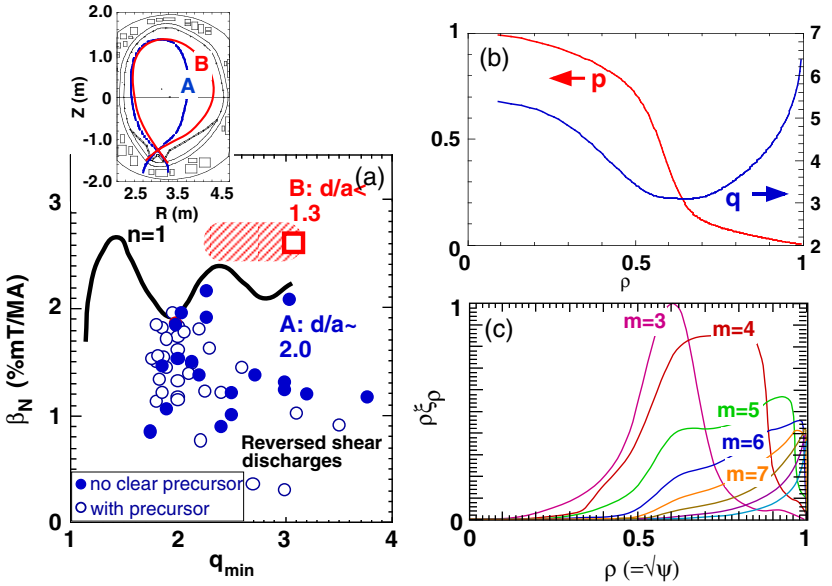

Figure 26. (*a*) Region of achieved β_N just before major collapse in the reversed shear discharges with $d/a \sim 2.0$ in JT-60U. The solid line indicates the computed stability boundary for the $n = 1$ kink-ballooning mode with the wall at infinity. An open square shows data with $d/a \sim 1.3$ and the hatched region indicate the limit in the discharges with $d/a < 1.3$. Radial profiles of (*b*) the pressure p and the safety factor q and (*c*) eigenfunctions of ideal $n = 1$ modes are for a discharge indicated as a square in (*a*). Reprinted with permission from [279].

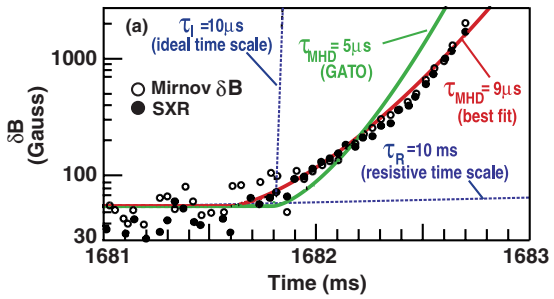

Figure 27. Time evolution of Mirnov δB (open circles) and SXR (filled circles) signals during the disruption in DIII-D showing the $\exp[(t/t)^{3/2}]$ dependence, from [259].

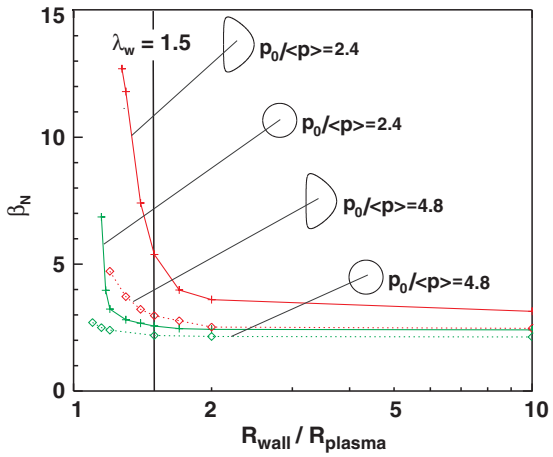

Figure 28. Dependence of the β_N limit on the conformal wall position $\lambda_w = R_{wall}/R_{plasma}$ for the four extreme cases with circular cross-section peaked (dotted line) and broad (solid line) pressure profile and similarly for the D-shaped with $\kappa = 1.8$ and $\delta = 0.7$. The scale for λ_w is logarithmic. The nominal wall position at $\lambda_w = 1.5$ used for most of the calculations is indicated by the vertical line [248].

then predicts quantitatively the non-exponential growth of the SXR and the Mirnov signals preceding the disruption, as shown in figure 27 [259]. The final disruption therefore appears to be quite consistent with the onset of an ideal instability, in terms of growth rate and mode structure. Further in ASDEX Upgrade, the ideal kink eigenfunction from the CASTOR code was found to be in good agreement with the experimental one derived from ECE just before the disruption [260].

Optimization of profiles of pressure, q, shaping and wall position is a key factor for improving the beta limit. Systematic calculations by varying the pressure peaking, the cross-section shape and the wall position showed clearly that there is a significant gain in all the figures of merit related to β, from the broadening of the pressure profile and strong shaping [248,262,263]. Stabilization by pressure broadening is similar to that in normal shear plasmas [264]; however, the gain in the β_N limit from a broad pressure profile was synergistic with additional large gains from strong shaping and wall stabilization. The effect from all three is multiplicative instead of additive. This is summarized in figure 28 showing the β_N limit as a function of the wall position for two extreme (circular and highly shaped) cross sections and two sets of pressure profile. Experimentally, an H-mode transition is used in DIII-D [262] and JT-60U [257] to broaden the pressure profile and this results in a significantly increased β_N limit, consistent with the systematic study and specific numerical calculation [265]. The effect of the edge current on global stability was also investigated for reversed shear discharges. The edge current density was shown to play an important role in coupling the internal ideal low n kink modes to the plasma edge, thus decreasing the ideal stability limit significantly [244]. The operational β_N limit in DIII-D has been found to increase with the shape factor from increased elongation and triangularity

for high-performance discharges with weak or intermediate shear [239, 254, 266]. This is consistent qualitatively with the systematic stability calculations [248]. Higher elongation is beneficial to all modes, while the combination of triangularity and outer squareness has a varying effect [267].

Localized minor collapses were observed near the transport barrier in JT-60U and TFTR [247] and $q > 1$ sawtooth-like events [258] in JET. The barrier localized mode (BLM) [244] was first observed in the high-β_p mode discharges in JT-60U, and the BLM was also observed in reversed shear discharges which was identified as an ideal MHD $n = 1$ instability [247]. The BLM does not degrade the improved confinement significantly but rather relaxes the strong ITB. Theoretical models are used to identify the dependence of the instability on the salient features of the plasma profiles; these are found to be the local pressure gradient, local shear in the q profile and the proximity of rational surfaces. The mode width is shown to depend on the value of q where the pressure gradient is largest—it increases as q decreases. Sawtooth-like events with $q > 1$ occur most commonly in the lower hybrid prelude phase, but can also persist into the main NBI/ICRF heating phase in JET [245, 258]. The sawteeth are associated with transient collapses of the ITB and are visible in the electron temperature during the LH prelude phase. In general, no precursor activity is observed before these sawteeth, but $n = 1$ post-cursors to the sawteeth are quite common and generally show no phase inversion of the electron temperature oscillations with radius, suggesting an ideal instability [258]. Other instabilities, including 'snakes' at the outer $q = 3$ surface, are also observed to limit the performance of reversed magnetic shear ITB regimes.

MHD modes such as double tearing modes (DTMs) or a mode coupling between two resistive modes on the inner and outer mode rational surfaces may also cause a major collapse [174,257,268]. On ASDEX Upgrade, a DTM occurring during the current ramp-up in strongly reversed shear scenarios led to a transient loss of the ITB. This could be avoided by additional ECR heating, which, according to modelling, prevents DTM formation due to increased conductivity [260].

Nonlinear extended MHD simulations were performed using the NFTC code [269] to resolve the cause and the role of the early MHD bursts. These found a linearly unstable resistive interchange mode at the inner resonant surface of the negative shear region and a double tearing mode, but the latter was stabilized by rotation. Without additional drive, the resistive interchange modes were found to saturate at low amplitudes, but they can be sustained in a metastable state.

Motional Stark effect measurements of the poloidal magnetic field in JT-60U [270] and JET [271] show that with large local bootstrap current or off-axis LHCD, the central current density reduces to near zero. Theoretical work has progressed on understanding the current hole equilibria. The standard equilibrium theory was re-examined and extended to include equilibria with a central current hole with no toroidal or poloidal current [272]. It was shown that equilibria with a central current hole can exist. One possible explanation for the equilibrium of a current hole is the 'axisymmetric tri-magnetic-islands (ATMI) equilibrium', as proposed in [273] which has three islands along the R direction (a central-negative current island and two side-positive-current islands) and two x-points along the Z-direction. The equilibrium is stabilized by the elongation coils when the current in the ATMI region is limited to be small. The idea was numerically confirmed in [274].

The linear analysis of MHD instability due to the negative toroidal current in the current hole shows that $m/n = 1/0$ resistive interchange modes with the growth rate $\sim \eta^{1/3}$ are unconditionally unstable, whereas $n = 0$, $m > 1$ tearing instabilities occur if $\Delta' > 0$ [275, 276]. Nonlinear simulations using the reduced MHD equations, including off-axis current drive, show that the $n = 0$ activity leads to cyclic reconnection events which clamp the central current density at approximately zero. Simulations in full toroidal geometry using the M3D code [277] also show a similar picture of cyclic $n = 0$ reconnection events, though these should be interpreted as being due to an axisymmetric equilibria not existing in this case [274]. As yet there is no direct observation of the MHD events causing the current clamping to zero. It is also important to clarify the beta limit for the strongly hollow current profiles due to the global MHD instability. The stability analysis of ideal MHD modes shows a possibility of significant improvement in the beta by adjusting the pressure profile [278].

2.5.3. Expectations for ITER. Ideal external kink modes coupled to 'infernal' modes due to the presence of low shear regions inside the plasma may limit the value of β_N in ITER advanced scenarios. The steady-state (SS) operation with $Q > 5$ requires a high $\beta_N > 2.5$, low internal inductance $l_i \sim 0.5$–0.7, high fraction of the bootstrap current (~50%), off-axis current density maximum and reversed magnetic shear (RS) [280]. The required $\beta_N > 2.5$ is close to or above the upper semi-empirical estimate of the no-wall stability limit $\beta_N^{\text{no-wall}}/l_i = 4$ [1] for this range of parameters (and slightly exceeds the calculated values), see table 3, section 2.3.3. Although the profile optimization for RS tokamak equilibria would allow for much higher values of $\beta_N^{\text{no-wall}}$ in ITER [248, 281, 282], realistic scenarios will restrict the accessible profile family. Therefore, resistive wall mode (RWM) stabilization is required (see section 2.3).

The SS scenarios in ITER were investigated in [203] with plasma transport simulated by the ASTRA transport code [283]. The ideal MHD stability analysis was performed for external kink modes using the KINX code assuming a separatrix at the plasma boundary [284]. The RS scenarios with different current profiles were provided by the variation of the lower hybrid (LH) current drive and the auxiliary heating and current drive (CD) by the neutral beam (NB) injection. The internal transport barrier (ITB) was modelled by a fast drop of transport coefficients to the level of the ion neoclassical diffusivity in the reversed shear zone. Since the neoclassical diffusivity in the RS zone increases with the safety factor, the increase in q_0 and $q_{\min}$ ($q_{\min} = 2.12$–2.43) reduces the pressure peaking factor (PPF) $p_0/ < p >$ (the ratio of central to average pressure) from 3.1 to 2.7 (figure 29). The reduction of PPF together with the deviation of $q_{\min}$ from 2 improves the stability, although it reduces the power multiplication factor Q. For the considered SS scenarios the possible operational limits were determined by rescaling the pressure profile keeping the current density and safety factor profiles unchanged. For such equilibrium series the values of the limiting wall position conformal to the plasma boundary for $n = 1$ mode were

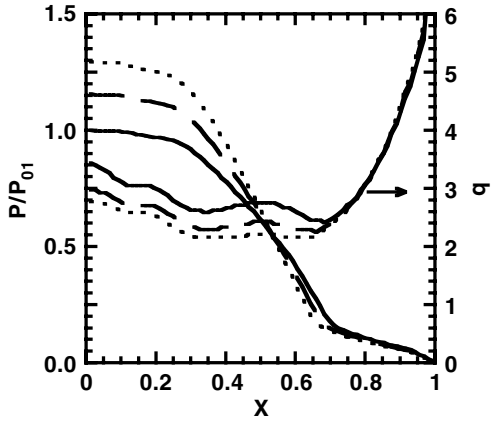

Figure 29. Normalized pressure and safety factor profiles for SS operational scenarios: solid line (N1): $q_{\min} = 2.43$, $\beta_N^{SS} = 2.56$, dashed line (N2): $q_{\min} = 2.25$, $\beta_N^{SS} = 2.7$, dotted line (N3): $q_{\min} = 2.12$, $\beta_N^{SS} = 2.82$; p_{01} corresponds to central pressure in the scenario N1 [203].

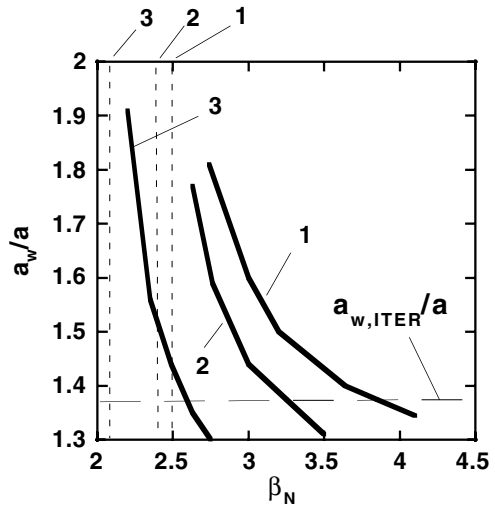

Figure 30. Stabilizing wall position a_w/a versus normalized beta β_N or q = const scan of SS operational points N1,2,3. The no-wall limits are shown by vertical dashed lines [286].

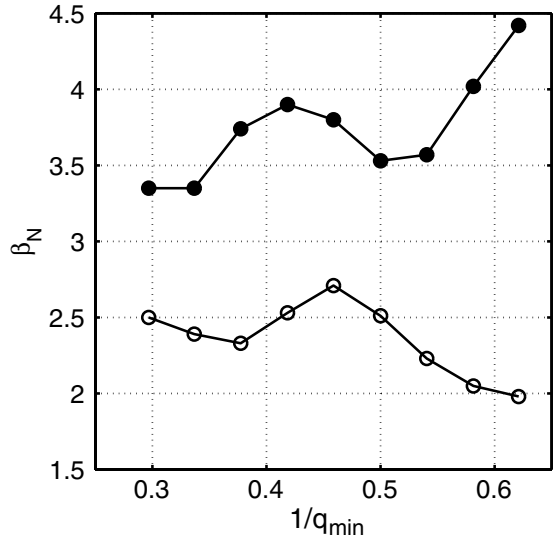

Figure 31. Limiting β_N versus the minimal safety factor value for the equilibrium series related to a full-bore scenario with $a = 2$ m. Open symbols correspond to $\beta_N^{\text{no-wall}}$, solid symbols to $\beta_N^{\text{ITER-wall}}$; $p_0/\langle p\rangle = 2.55$, $l_i = 0.58$.

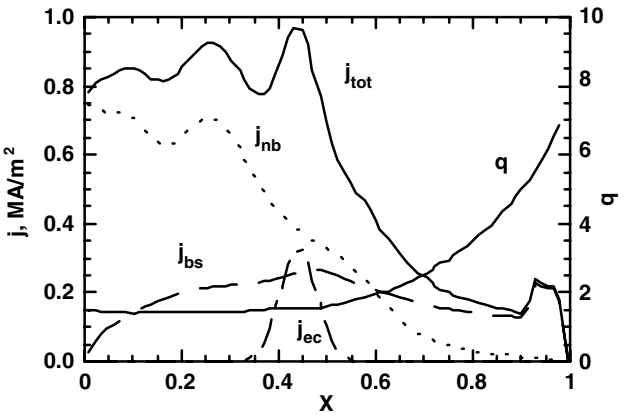

Figure 32. Current density and safety factor profiles for the SS equilibrium with $\beta_N^{SS} = 2.67$, $q_0/q_{\min}/q_{95} = 1.6/1.45/5.25$, $p_0/\langle p\rangle = 3.5$, $l_i = 0.8$.

calculated (figure 30). The stability calculations showed that the stabilizing influence of the ITER vacuum vessel on the n=1 mode is well described by the conformal wall with $a_w^{ITER}/a = 1.375$ for plasma minor radius $a = 1.85$ m (plasma centre $R = 6.35$ m).

A full-bore scenario with $Q = 8$ was also considered ($a = 2$ m, $R = 6.2$ m, $a_w^{ITER}/a = 1.35$) with higher $\beta_N^{SS} = 3.6$ at the SS operational point. Due the smaller the plasma–wall distance and the lower value of PPF $p_0/ < p >= 2.55$ ($q_{\min} = 2.16$) the value of $\beta_N^{\text{ITER-wall}} = 3.8$ is also higher in this case ($\beta_N^{\text{no-wall}} = 2.7$). It gives quite a high value of $C_\beta = (\beta_N - \beta_N^{\text{no-wall}})/(\beta_N^{\text{ITER-wall}} - \beta_N^{\text{no-wall}})$, which determines a degree of RWM instability (see section 2.3.3): $C_\beta = 0.82$ for $\beta_N^{SS} = 3.6$. The stability study for different values of normalized current I_N was performed in [207] using the profiles from the transport calculations. The results in the (β_N, $1/q_{\min}$) parameter plane (figure 31, $I_N \approx 2.5/q_{\min}$) show that despite an overall decrease in the $\beta_N^{\text{no-wall}}$ values along with lower $q_{\min}$, the wall stabilization becomes stronger for higher I_N. This can be explained by the fact that the mode structure changes with increasing I_N: the mode becomes more extended to the low-field-side and couples more strongly to the wall. For the lowest value of $q_{\min} = 1.6$ the increase in $\beta_N^{\text{ITER-wall}} = 4.4$ ($\beta_N^{\text{no-wall}} = 2.0$) corresponds to $C_\beta = 0.67$ provided that the steady-state value β_N^{SS} is preserved.

Recent modelling of the SS scenarios with higher internal inductance $l_i \approx 0.8$ and lower $q_{\min} \approx 1.5$ resulted in safety factor profiles with weakly reversed shear in the plasma center [285] (figure 32). As expected, it gives an increase in the β_N limits against an external n=1 mode ($\beta_N^{\text{no-wall}} \approx 3$), but this configuration is unstable against the $n = 2$ mode once the resonant surface $q = 1.5$ is in plasma. The $n = 2$ RWM stabilization is an issue and NTM mode stabilization can be also required for such profiles.

In summary, the SS operation in ITER requires RWM stabilization for the scenarios with $q_{\min} > 2$. The SS operational point with relatively low pressure peaking

$p_0/ < p > \approx 2.5$ should be considered to provide a wider range of permitted pressure excursions $\beta_N^{\text{ITER-wall}}/\beta_N^{\text{SS}} \approx 1.5$. In the scenarios with higher l_i, the values of $\beta_N^{\text{no-wall}}/\beta_N^{\text{SS}} \approx 1.2$ can be expected even with higher PPF and without wall stabilization provided that $q_{\min} > 1.5$. In both cases, sufficiently accurate profile control would be desirable to avoid shrinking of the operational space.

2.6. Summary

Recent theoretical and experimental research has made important advances in both understanding and control of MHD stability in tokamak plasmas. Central $m/n = 1/1$ sawteeth are anticipated in ITER, but the tools exist to avoid or control this instability through localized current drive or fast ion generation. Active control of other MHD instabilities will most likely be required in ITER. Extrapolation from existing experiments indicates that stabilization of neoclassical tearing modes by highly localized, feedback-controlled current drive should be possible in ITER. Resistive wall modes are a key issue for advanced scenarios, but again extrapolation from existing experiments indicates that these modes can be stabilized by a combination of plasma rotation and direct feedback control with non-axisymmetric coils. Reduction of error fields is a requirement for avoiding non-rotating magnetic island formation and for maintaining plasma rotation to help stabilize resistive wall modes. Recent experiments have shown the feasibility of reducing error fields to an acceptable level by means of non-axisymmetric coils, possibly controlled by feedback. The MHD stability limits associated with advanced scenarios are becoming well understood theoretically and can be extended by the tailoring of the pressure and current density profiles as well as by the other techniques mentioned here.

The $m = 1$ internal kink instability is an important issue for present tokamak experiments and future burning plasmas. A large sawtooth crash can trigger MHD instabilities such as the neoclassical tearing mode, leading to confinement degradation, locked modes and/or disruptions. Therefore, it is important to develop methods for controlling sawteeth in present and future experiments. A wide range of techniques have been demonstrated for eliminating sawteeth or reducing their impact, including stabilization by fast particles, modification of the current density profile with localized current drive or stimulation of small amplitude sawteeth with local heating or current drive. Neutral beam heating, ion cyclotron heating, electron cyclotron current drive and lower hybrid current drive have all been used for this purpose. Of particular interest to burning plasmas is the effect of a significant population of fast alpha particles on the stability properties of the $m = 1$ mode and the redistribution of the alpha population following a sawtooth crash. Although a detailed physics understanding of the $m = 1$ mode remains an open subject, a predictive capability is emerging in the form of models for the triggering of $m = 1$ instabilities that are based on MHD stability with kinetic effects of thermal and fast ion populations. These models are coupled with transport codes to predict the effects of sawteeth on the total stored energy and fusion power. The results, benchmarked to experimental data, suggest that in a burning plasma such as ITER sawteeth have little effect on the total stored energy or time-averaged fusion power.

Neoclassical tearing modes represent a significant challenge for present tokamaks and future burning plasmas, but the tools now appear to be at hand to control or suppress these instabilities. Recent experimental studies and simple scaling models based on modified Rutherford theory indicate that the critical normalized beta above which NTMs can potentially be destabilized scales approximately linearly with the normalized ion gyroradius $\rho_i^* = \rho_i/a$. This scaling, although unfavourable for large burning plasma experiments, does not directly imply lower NTM beta limits, since the scaling and nature of the NTM seeding mechanisms is the key factor and this issue is not yet fully resolved. In any event, recent experiments have been quite successful in using localized electron cyclotron current drive to stabilize the $m/n = 3/2$ and 2/1 NTMs, allowing beta to be raised above the threshold value observed in the absence of current drive. Other promising techniques for avoidance or mitigation of NTMs include the control of sawteeth (and hence the generation of seed islands), the control of the plasma shape and profiles, elimination of low-order rational surfaces (e.g. 3/2 and 2/1) from the plasma and operation in the 'frequently interrupted regime' where nonlinear coupling by other MHD modes limits the amplitude of the longer wavelength, more deleterious NTMs. Given an appropriate launcher and sufficient power for electron cyclotron current drive, it should be possible in ITER to keep NTMs at a small amplitude or to stabilize them completely.

Advanced scenarios in present tokamaks and in ITER often rely on wall stabilization of the $n = 1$ kink mode at high beta. In the presence of a wall with finite conductivity, an external kink that would be stabilized by a perfectly conducting wall instead becomes a slow-growing resistive wall mode. In present experiments the resistive wall mode can often be stabilized by plasma flow (at resonant surfaces such as $q = 2$ or $q = 3$) on the order of a per cent of the Alfvén speed or less [154, 155]. This result is in reasonable agreement with theoretical predictions, although precise predictions are sensitive to the plasma profiles and the choice of model for dissipation in the plasma. Optimized correction of magnetic field errors can help in maintaining the plasma rotation. However, a burning plasma will have little external torque applied to it, and the rotation rate is likely to be marginal at best. Recent experiments have shown success in active feedback control of the resistive wall mode, using non-axisymmetric coils both external and internal to the resistive wall, and also show that a sub threshold rotation can be combined with feedback stabilization for increased stability. Modelling predicts, and early experimental results confirm, that internal coils are capable of stabilization at higher beta. Modelling indicates that the modest improvement in beta over the no-wall limit needed for advanced scenarios in ITER can be realized with feedback stabilization using the external error correction coils. However, a significant reduction in the current and the voltage of the control coils, as well as a greater improvement in beta, could be achieved with internal coils.

Recent theories and experiments have emphasized the need to reduce magnetic field errors. Relatively small non-axisymmetric magnetic fields having components that are resonant with low-order rational surfaces in the plasma can slow plasma rotation, cause locked modes or generate seed

islands for neoclassical tearing mode growth. Slowly rotating plasmas are particularly susceptible to the penetration of field errors. Present experiments have been quite successful at reducing the effects of error fields by the use of non-axisymmetric correction coils, and such coils have become a feature of most major tokamaks. Analytic estimates combined with scaling from present experiments give a low but achievable level of an acceptable error field for ITER. The error field limits become somewhat smaller for advanced scenarios in order to avoid enhanced rotational drag from the 'error field amplification effect,' a resonant response by marginally stable resistive wall modes in plasmas above the no-wall stability limit. However, reduction of error fields to a tolerable level should still be achievable with the correction coil systems planned for ITER.

The special features of advanced scenarios present specific challenges to MHD stability, some of which have been mentioned above. Non-inductive operation without large recirculating power requires a large bootstrap current, which in turn requires operation at lower plasma currents and higher normalized beta. The bootstrap current tends to produce a broad current density profile with elevated central safety factor, leading to small or negative magnetic shear in the core of the plasma. The broad current density profile requires wall stabilization of the external kink mode at high beta, as discussed above. A central region of negative magnetic shear is predicted to be stable to neoclassical tearing modes, but may be unstable to resistive interchange modes; experimental results indicate that the interchange modes may place limits on the central peaking of pressure, the rotation and the safety factor. Double tearing modes are also a potential problem in discharges with negative central shear but are probably stabilized by rotational shear in most present experiments. A more serious issue is the ideal kink or the 'infernal' mode which may be destabilized in a central region of weak shear and strong pressure gradient; internal transport barriers can contribute to strong local pressure gradients. Several experiments have observed disruptions resulting from such instabilities, and experiment and theory show that the beta limit varies inversely with the peaking of the pressure profile. The experiments also show that such disruptions can be avoided through the control of the pressure profile, by means of off-axis heating or by profile broadening with a controlled H-mode transition. Accordingly, it is expected that these instabilities can be avoided in ITER through wall stabilization and control of the pressure profile peaking.

In summary, recent research has deepened our understanding of MHD stability in tokamaks and has demonstrated approaches to avoidance or direct control of instabilities. Although MHD stability remains an important issue for future experiments such as ITER, we can say with some confidence that MHD stability limits will not be a fundamental obstacle.

3. Disruptions

Disruptions and their consequences pose significant design and plasma operation challenges for reactor-regime tokamaks in general and for ITER in particular. The magnitude and scope of these challenges arise from a combination of physics, structural and thermal engineering considerations and from inherent limits on the thermal energy handling capabilities of materials available for plasma-facing component (PFC) surfaces. On the structural engineering side, a burning-plasma-capable tokamak based on niobium–tin superconducting magnet technology allows a plasma-axis toroidal magnetic field of about 6 T, an increase in field relative to the 3–4 T fields routinely employed in the present generation of large, high-performance tokamaks. Further the associated increase in plasma current to about 15 MA brings with it a moderate ($\sim$3) increase in the electromagnetic (EM) $B_{\rm pol}^2/2\mu_{\rm o}$ pressure loadings on the torus vacuum vessel (the comparison here is with the corresponding magnetic pressure loadings in the present generation of large tokamaks, see table 5). While the structural loading implications of these higher EM forces are non-trivial, acceptable engineering design solutions have been identified for past [287] and present [288] ITER designs.

Disruption-produced thermal energy loading on the PFC surfaces of an ITER-class tokamak poses a less easily solved challenge [287]. Here the predicted time-normalized surface energy loading ($U/(A \times t^{0.5})$, where U is the deposited energy, A is the area of deposit and t is the deposition time) on the divertor targets will be high enough that some localized melting and/or vaporization of the affected surfaces (typically either W or C) will be likely. While the time-normalized energy loadings for the present 15 MA ITER design concept are appreciably lower than those predicted on the same basis for the 21 MA EDA design (see table 5), and while, in any case, the depth of the melt- or vaporization-affected layer produced by a single disruption will be small ($\leqslant\sim$100 μm), the erosive effects of a series of unmitigated disruptions in the present ITER would still likely impact the useable lifetime of the divertor PFC surfaces. Furthermore, the predicted sublimation of carbon following a disruption raises serious plasma operation concerns about de-conditioning of the divertor and the torus PFC surfaces and tritium retention in co-deposited carbon layers. Accordingly, there will be a strong operational incentive to both reduce the frequency of occurrence of disruption and (at very least) reduce or ameliorate the thermal loading consequences of disruptions that cannot otherwise be avoided.

The production during disruptions of relativistic (runaway) electrons poses a second type of threat to the integrity of the ITER PFC surfaces. ITER, like any high-current reactor-regime tokamak, will be inherently susceptible to efficient conversion, by Coulomb-collision avalanche multiplication, of plasma current to relativistic (runaway) electron current [289,290]. Such conversion, of up to 70% of the initial plasma current, is predicted to occur following either a naturally occurring disruption or an artificially induced fast plasma shutdown. The subsequent uncontrolled interaction of this magnitude of runaway current with PFC surfaces has the potential to produce local damage to PFC surfaces and their underlying substrate structures. Accordingly, ITER (and future reactor tokamaks) will almost certainly need to have disruption mitigation and fast-shutdown means that can simultaneously ameliorate disruption thermal loading and runaway conversion.

Table 5 compares various disruption-related parameters for JET (chosen here as a representative 'large-tokamak' example), ITER [288] and the 21 MA ITER-EDA [287] designs. The data in Table 5 are presented here to give a

Table 5. Disruption and disruption consequences for JET, ITER and ITER-EDA.

Parameter	JET	ITER	ITER-EDA	Basis or comment
R (m)	2.9	6.2	8.14	Major radius
a (m)	0.95	2.0	2.8	Minor radius
κ_{95}	1.6	1.7	1.6	Vertical elongation
V (m^3)	86	831	2000	Plasma volume
S (m^2)	145	683	1200	Plasma surface area
B_T (T)	3.45	5.35	5.68	Toroidal field
I_p (MA)	4.0	15	21	Plasma current
q_{95}	3.0	3.0	3.0	Edge safety factor
W_{mag} (MJ)	~11	395	1100	Poloidal field energy inside separatrix
W_{th} (MJ)	~12	353	1070	$\beta_N = 2$, with 'ITER-like' $p(r)$ profiles
Magnetic and current quench related attributes				
$\langle B_p \rangle$ (T)	0.60	1.07	1.13	Average poloidal field
$\langle B_p \rangle^2/2\mu_o$ (MPa)	0.143	0.454	0.507	Torus vacuum vessel magnetic pressure
t_{CQ} (ms)	9.4	35.6	65.7	Minimum current quench duration
$B_T * dB_p/dt$ ($T^2\,s^{-1}$)	220	161	98	Relative force due to induced eddy currents
$W_{mag}/(A_{FW} * t_{CQ}^{0.5})$ (MJ $m^{-2}\,s^{-0.5}$)	0.78	3.1	3.6	cf Be melt onset at ~15 MJ $m^{-2}\,s^{-0.5}$
I_{halo}/I_p	⩽0.45 (data)	⩽0.4 (est)	⩽0.4 (est)	Halo current fraction
TPF	⩽1.7 (data)	⩽2 (est)	⩽2 (est)	Toroidal peaking factor
Thermal quench and divertor energy loading attributes				
A_{div} (m^2)	~1.6[a]	~3.5	~4.6[a]	Effective divertor target area, for H-mode
$U_{TQ} = W_{th}/7A_{div}$ (MJ m^{-2})	1.07	14.1	33	For 7-x SOL expansion during-disruption TQ
t_{TQ} (ms)	0.32	0.70	1.0	As per figure 54 of [1]
$U_{TQ}/t_{TQ}^{(0.5)}$ (MJ m^{-2})	60	530	1040	C or W vapour/melt onset at 40–60 MJ $m^{-2}\,s^{-0.5}$
Runaway electron conversion and mitigation attributes				
E_{int} (V m^{-1})	38.3	38	28.8	In-plasma E-field
$n_{e,RB}$ (m^{-3})	4.2×10^{22}	4.2×10^{22}	3.2×10^{22}	n_e to suppress avalanche growth
$G_{avalanche}$	2.2×10^4	1.9×10^{16}	6×10^{22}	Coulomb avalanche gain = exp[2.5 × I (MA)]
$I_{RA,\ seed}$ (A)	90	4×10^{-10}	1.8×10^{-16}	Seed current for $I_{RA} = 0.5I_p$
t_{fs} (ms)	0.030	1.2	3.5	Minimum W_{th} shutdown time to avoid Be FW melt

[a] Divertor area estimates for JET and ITER-EDA assume an R^1 scaling of A_{div}.

comparative assessment of the relative 'challenge' of the three major ITER disruption consequence issues—EM loading, thermal loading and runaway conversion. The physics basis considerations that enter into table 5 have been extensively detailed in chapter 3, section 4 (disruption and disruption effects) of the ITER Physics Basis (IPB) [1]. Further discussion of the present status of this physics basis and of new interpretations relative to what is detailed in the IPB follow in sections 3.1–3.6.

Table 5 demonstrates several of the points noted above: EM pressures on the ITER torus vacuum vessel associated with the plasma current quench are about three times higher than the corresponding pressures in JET (if JET had an ITER-like low-resistance vacuum vessel), and the local induced eddy current forces on ITER first wall and RF antenna structures will be somewhat lower than the corresponding forces on JET structures (of otherwise similar construction). Finally, the thermal loading on the ITER first wall due to radiative dissipation of the plasma magnetic energy during the current quench phase is well below the $W/(A_{FW} \times t_{CQ}^{0.5})$ threshold (~15 MJ $m^{-2}\,s^{-0.5}$) that applies to the onset of surface melting of beryllium. Accordingly, the conclusion originally given in the ITER Physics Basis [1] that the EM loading and structural engineering and first-wall PFC challenges, posed by the fastest-expected ITER current quench can be successfully accommodated will continue to apply for the present 15 MA ITER design. Note, however, that the eddy current-induced loadings in ITER are now somewhat higher than those that are applied for the ITER-EDA. Hence, there is renewed interest in understanding the bounds on the maximum current quench rate and the corresponding dB_p/dt expected in the present ITER design. New findings in this regard are described in section 3.3.1.

The present status of the physics basis for VDEs and for estimating halo current magnitude and toroidal asymmetries (typically described in terms of the toroidal peaking factor or TPF) and the resulting vacuum vessel forces due to VDEs are described in section 3.3.2. The general conclusion reached here and in [288] is that there has been little change in the physics basis status and the recommendations for ITER design basis guidelines on the maximum halo current magnitude $(I_{halo,max}/I_{p0})$, TPF and the product $(I_{halo,max}/I_{p0}) \times TPF$.

The third current quench phase related subject, that of runaway electron generation, amplification and loss (to PFC surfaces) is addressed in section 3.4. As the section details, in a high-current tokamak such as ITER, the effect of Coulomb avalanche multiplication is predicted to provide a much stronger coupling between the effects of toroidal plasma current decay and equilibrium dynamics (VDE) and runaway production and ultimate loss to PFC surfaces. The surface-damage potential of interaction of a multi-MA runaway current with localized portions of the at-risk PFC surfaces leads to serious concerns about the high levels of runaway conversion following naturally occurring disruption and intentional 'fast-shutdown' actions, intended to ameliorate disruptions.

The three current quench related sections (3.3.1, 3.3.2 and 3.4) are preceded by two preliminary sections. Section 3.1

addresses both the overt causes for disruption and the resulting internal MHD reconnection and flux surface destruction that leads to thermal quench, current quench and VDE 'consequences' of disruption. The subject of disruption frequency and 'causes', operational and otherwise, is also addressed in this section.

Section 3.2 addresses the physics basis and consequences of the 'thermal quench' phase of a disruption. Here thermal quench denotes the rapid internal redistribution and subsequent loss to PFC surfaces of the plasma thermal energy, $W_{\rm th}$, that occurs immediately before the onset of disruption. The well-known feature of this phase is the rapidity of the internal redistribution and/or thermal energy loss, in times that are predicted to be as short as ~1 ms for an ITER-class tokamak. The thermal-quench portion of table 5 summarizes some of the key attributes for present tokamaks and ITER designs. The table demonstrates that while time-normalized thermal quench loadings for ITER are now significantly reduced relative to those projected on the same basis for ITER-EDA, it is still likely that some divertor PFC surface melting or vaporization might occur. However, the depth of the melt- or vaporization-affected zone will be reduced relative to the depth projected for ITER-EDA, and hence the cumulative erosive effects of a given frequency of unmitigated disruptions may be less PFC-lifetime limiting than was projected—on a worst-case basis—for the ITER-EDA. Furthermore, new physics basis details on the time dynamics and spatial partitioning (among divertor and other PFC surfaces) and the spatial spreading of plasma thermal energy deposition on the divertor surfaces during disruption suggest that melt- or vaporization-free divertor PFC operation may be obtainable in at least some of the proposed ITER plasma operation regimes. These important new findings about thermal quench characteristics and discussion of the resulting implications for the ITER divertor PFC response are given in section 3.2.

Section 3.5 discusses recent advances—in most cases, using some degree of an integrated model—in the numerical simulation of ITER disruptions and VDEs, and their resulting EM loading and runaway electron generation consequences. Representative examples that contribute to reaching the conclusions noted above about the structural integrity of ITER vacuum vessel and PFC substrate systems are presented. Progress in incorporating more-fundamental MHD instability considerations into these types of simulations is also addressed.

Finally, section 3.6 takes up the closely-connected topics of disruption prediction, avoidance, and mitigation means and their projection to ITER. The importance of suppressing runaway avalanching underlies such projection and sets rather stringent limits on whatever technique is to be implemented.

Section 3.7 provides a summary of the present understanding of the physics basis for disruption and disruption avoidance and mitigation in ITER and gives recommendations as to needs for future physics and technology R&D.

3.1. Disruption characteristics, causes and frequency

It is well known that stable sustainable (disruption-free) operation in a tokamak system is limited with regard to maximum plasma current, maximum electron density and maximum total normalized plasma pressure (β) by three basic 'operational limit' considerations:

- Current 'limit': set by a requirement for a plasma edge safety factor, $q_{95} \geqslant\sim 2$.
- Density 'limit': set by a requirement that the plasma density should not appreciably exceed the empirical 'Greenwald' density limit $n_{GW}(10^{20}\,\mathrm{m}^{-3}) = I\,(\mathrm{MA})/\pi a^2(\mathrm{m}^2)$.
- Pressure 'limit': set by a requirement that the normalized volume-average toroidal beta, $\beta_{\rm N} = \langle\beta_{\%}\rangle/(I\,(\mathrm{MA})/a(\mathrm{m})\,B(\mathrm{T}))$, should not exceed the 'Troyon' ideal MHD beta-limit of approximately 3.5% MA m^{-1} T^{-1}.

The plasma current and pressure limits reflect the immediate effects of onset of ideal MHD instability, whereas the 'Greenwald limit' manifestation of the plasma density limit, which appears to be the result, in part, of an underlying deterioration of plasma energy and particle confinement, may have a more indirect MHD origin that in turn leads to an MHD-initiated disruption. The subject of plasma density limits has been discussed in detail in chapter 3 of IPB [1] and also, with regard to the role of intermediate MHD instability, in the material that follows below.

Attempting plasma operation that approaches or exceeds the limiting values of any of these three operational boundary 'limits' typically initiates an increase in MHD activity that eventually results in a major disruption, wherein the tokamak magnetic configuration becomes globally unstable to helical perturbations of the form $\zeta(r,t) = \zeta(r)\exp \mathrm{i}(m\theta - n\phi)$, where θ and ϕ are the poloidal and toroidal angles and m and n are the corresponding poloidal and toroidal integer mode numbers. Once the final non-linear growth phase of this helical instability begins, global destruction of the nominal integrity of the nested flux surfaces develops on a time scale that can be as short as 100 μs in the present generation of medium-size tokamaks. Extrapolation (see chapter 3, section 4 of [1]) to ITER predicts a corresponding time scale for thermal energy redistribution within the plasma volume that is ~1 ms.

As a result of this rapid growth of global MHD instability, major disruptions in present tokamaks can expel most of the plasma thermal energy on a time scale that is also typically ~100 μs. The ensuing increase in plasma resistivity that this 'thermal quench' cooling produces then precipitates a rapid decay of the plasma current (typically described as a 'current quench') and, in a vertically elongated tokamak, simultaneous development of vertical instability, typically described as vertical displacement event (VDE) and sometimes also as a vertical disruption or vertically unstable disruption. Current quench and VDE development occur on a time scale that is typically ~5 ms in present medium-size tokamaks. The corresponding current quench time scale extrapolated to ITER is about 35 ms.

The rapid current quench that follows disruption also gives rise to a high in-plasma electric field that can generate superthermal (runaway) electrons, with energies in the 10–100 MeV range. In some cases, conversion of up to about 50% of the before-disruption plasma current to runaway current is observed in present 'large' tokamaks (see, e.g. [291] and section 3.4). Once created, the subsequent interaction of such a 'runaway discharge' with the background thermal plasma is weak and the runaways typically remain well confined for

as long as the plasma equilibrium is maintained. The eventual decay of the plasma equilibrium leads to impact of the runaway electrons on the PFC surfaces. Such impact can cause localized surface and substrate (below-surface) heating, vaporization and thermal shock damage and, in metals, recrystallization and crack development in the heat-affected zone.

The three operational limits ('boundaries') cited above as defining the disruption-free operation domain for tokamak plasmas are indicative, but not definitive. This means that while approaching or crossing one or more of these boundaries will usually lead to onset of disruption, there do exist examples where one or more of these 'conventional' operational boundaries can apparently be exceeded, and conversely, other examples where disruption occurs under conditions remote from all three boundaries e.g. disruptions initiated by tearing modes in high-q Ohmic 'start-up' plasmas (see chapter 8, section 3.2 of [1]) or disruption initiated by neoclassical tearing mode (NTM) onset in moderate-β_N positive-shear H-mode plasmas (see section 2.2), or by onset of pressure-gradient-driven modes in high-q 'advanced-performance' plasmas with internal transport barriers (see section 2.5 and e.g. [245]). The complexity of the underlying MHD stability and the possibility for 'secondary' energy confinement and/or MHD instabilities to in turn initiate the onset of primary ideal-MHD instability and disruption, make disruption prediction a much more challenging matter than simply evaluating proximity of the plasma attributes to the three 'limit' parameters (see section 3.6).

Mechanisms for major disruption. A major disruption is similar to a large-scale interchange-type MHD instability that simultaneously envelops both the plasma centre and its periphery. This type of interchange can be explained in two ways: magnetic reconnection [292] and non-linear ideal-helical instability (cold bubble) in regions of low shear [293]. Magnetic reconnection of plasma areas with different helicity can occur during the development of non-linear kink and tearing modes ($m = 1/n = 1$, $m = 2/n = 1$). These modes result in convective and conductive plasma transport, which leads to a rapid loss of thermal and magnetic energy in the central area of the plasma column. This internal disruption is similar to a sawtooth oscillation (see section 2.1), but it modifies the current distribution and magnetic shear not only close to the radius of inversion ($q(r) = 1$), but over most (or all) of the plasma radius.

In the 'cold bubble' model, helical flux tubes can be captured in the periphery and transported into the hot centre. This interchange ultimately leads to an increase in global resistivity and to the plasma current quench. The experimentally observed fast impurity transport during sawteeth [294] is in agreement with this interchange-type mechanism.

The thermal quench that this mixing precipitates is typically followed by a redistribution of toroidal current, leading to a flattened current profile (and hence to a decrease in the dimensionless plasma internal inductance, l_i). In order to conserve flux [295], the plasma current increases, producing a characteristic current overshoot or 'spike' ($\sim$10% I_{p0} for conventional aspect ratio tokamaks) and a large transient negative toroidal loop voltage. This transient voltage is

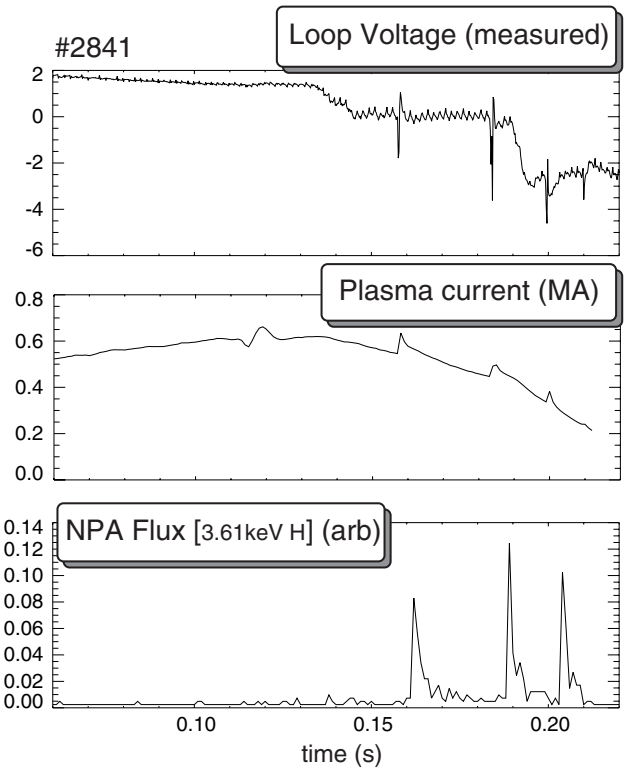

Figure 33. Edge loop voltage, plasma current and NPA flux at 3.61 keV in a MAST discharge with internal reconnection events. (Reprinted figure with permission from [296]. Copyright 2002 by the American Physical Society.)

readily observed at the plasma edge. The presence of such a large toroidal voltage at the plasma centre has been recently confirmed with non-magnetic evidence: figure 33 shows evidence for the toroidal acceleration of protons in the centre of a MAST plasma during internal reconnection events (partial disruptions) [296].

A detailed account of the dynamics of major disruptions for plasmas close to the ideal MHD beta-limit can be found in [297]. These dynamics include the development of NTMs and ballooning modes, followed by a deep 'internal disruption' and finally by a conventional major disruption (the typical two-step major disruption model cited in chapter 3, section 4 of [1]). The outer shell of hot dense plasma between the plasma edge and the zone of development of the sawtooth internal disruption may play a role in screening interaction with external $m = 2$–4, $n = 1$ modes [298]. A numerical modelling effort devoted to investigate the non-linear evolution of ballooning modes in high-temperature plasmas elucidates some of the MHD mixing mechanisms believed to be present. The numerical modelling shows that convection cells develop between the hot central plasma and the edge in ridges whose two-dimensional poloidal projection resembles fingers [299]. As the plasma resistivity decreases (is reduced as a parameter in the simulation), the number of fingers increases, while the width of each individual finger becomes narrower (figure 34). Because of these effects, the rate at which energy is transported to the wall (inferred from the simulation) is unchanged as the resistivity is reduced. The magnetic field lines are decoupled from the plasma even when the resistivity approaches zero,

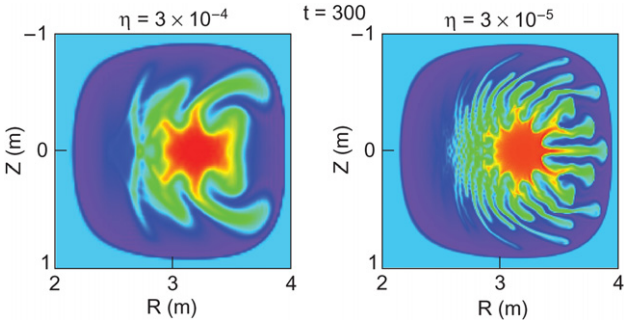

Figure 34. Temporal evolution of the pressure in the poloidal plane at a fixed toroidal angle. The normalized resistivity used in the simulations is 3×10^{-4} (left) and 3×10^{-5} (right). (Reprinted with permission from [299]. Copyright 2001, American Institute of Physics.)

resulting in the plasma fingers being readily convected across the magnetic field. Disruptions at high-β have also been simulated using the NIMROD code [300].

This prediction of fast mixing even for a low-resistivity plasma is in contrast to the predictions of disruption models based on magnetic stochasticity and/or magnetic reconnection. In resistive MHD, the rate of reconnection of magnetic field lines scales with the square root of the local resistivity (η) and there is a large reduction in reconnection rate for η magnitudes that are characteristic of the initial (before-thermal quench) phase of disruption. Therefore, reconnection models tend to predict reconnection times that increase with increased plasma temperature and/or reduced η. In contrast, nonlinear simulations demonstrate that confinement is destroyed by ballooning modes at a rate that is independent of the η and on time scales that are commensurate with the sudden loss of thermal energy observed during β-limit disruptions. Although these results strictly apply only to ballooning modes in a tokamak, these results may also be relevant to other MHD instabilities involving convection cells [301]. Their inclusion in an 'integrated' numerical simulation model (see section 3.5) may ultimately provide a first-principle basis for modelling the onset and thermal quench phase of a major disruption.

A sudden internal degradation of the energy confinement—a frequently observed feature of 'density limit' disruptions (e.g. a disruption caused by proximity to the Greenwald limit and/or by an excessive plasma edge fuelling rate)—has now been observed to initiate close to the $q = 2$ surface in small [302], medium [303] and large-sized [304] tokamaks. In each instance, a flattening of the electron temperature profile spreads inwards from the outboard (large-R) O-point of the 2/1 mode. The plasma periphery cooling that this spreading causes initiates destruction of the global energy confinement. The observation of the same onset phenomena in three distinct-size tokamak plasmas indicates that it may be an intrinsic characteristic of this type of disruption (perhaps with an underlying MHD cause). An increase in density fluctuation levels and an abrupt increase in the frequency of fluctuations were also observed to coincide with the onset of the electron temperature erosion. There are future prospects for incorporating such considerations in models of the precursor and onset phases of the major disruption.

Disruption causes and frequency. As the discussion presented above, and preceding discussions developed in the IPB make clear, onset of rapidly growing global MHD instability is always the penultimate cause of a major disruption. However, from a plasma operation point of view, there is also a proximate and usually clearly identifiable precursor 'cause' for each disruption that occurs. Such causes can be categorized either in terms of (1) the precursor MHD instability or plasma energy confinement or energy balance disturbance event that triggers the final global instability onset or (2) the operational or tokamak system event that 'causes' the disruption. Categorization in terms of item (1) leads to terminologies that include 'density limit' disruptions, 'cold-edge' disruptions, 'beta-limit' disruptions (either from ideal MHD or, more frequently in present experiments, from NTM or tearing mode growth and/or mode locking), and 'internal pressure-gradient' (ITB) triggered disruptions. Since the time scale for confinement deterioration can approach the (degraded) energy confinement time scale, the pre-MHD-onset phase of these types of events are often described as 'thermal collapse'. Categorization in terms of item (2) leads to a long list of hardware and/or operation-associated causes that can include hardware 'failures', e.g. premature turn-off of the Ohmic drive power supply or neutral beam heating, excessive gas fuelling input rates, release of impurities from PFCs owing to excessive power loading, debris falling from in-torus surfaces into the plasma, poor wall conditions (perhaps as the result of a previous disruption), plasma control system failure or inadvertent mis-programming (i.e. human error), intentional 'killer pellet' or massive gas injection, and so forth. These lists can be quite detailed and hardware specific and are not necessary universal, in terms of categories and terminology, across the present spectrum of operating tokamaks.

Most presently operating tokamaks now keep shot logs that describe both disruption 'type' or physics cause(s) and also the underlying operations-related event that the operator may identify as being the 'cause' of the disruption. These logbook entries constitute a 'database' that can be used to compile 'statistics' that quantify disruption-type, causes (physics, operation intent and/or hardware and human factors), development sequences (e.g. ne-limit → thermal collapse → disruption) and frequencies of global and/or by-type or by-category occurrence. Interpretation of such log-book data to search for physics-based correlations for disruption likelihood ('disruptivity') can, however, prove problematical. Figure 35 shows one such long-term assessment for DIII-D: there is little evidence for any strong correlation of long-term average disruptivity with any of the traditional 'operation limit' plasma parameters [305]. The three-year average disruptivity (with obvious hardware failure causes screened out) is about 13%. Statistical assessments of the same ~3600 discharge data set shows that the per-unit-time disruptivity of long-pulse high-β_N discharges, that successfully reach a stable stationary condition after about 4 s (from heating initiation), tends, within the statistical accuracy possible in the 300-discharge dataset, to zero as the discharge duration increases towards 7 s (set by hardware limits). This per-unit-time disruptivity decrease rules out, on a statistical basis, that disruptions occurring during this 4–7 s interval are precipitated by a random process.

Interpretation of these types of statistics on disruption causes and occurrence frequency continues to be applied to

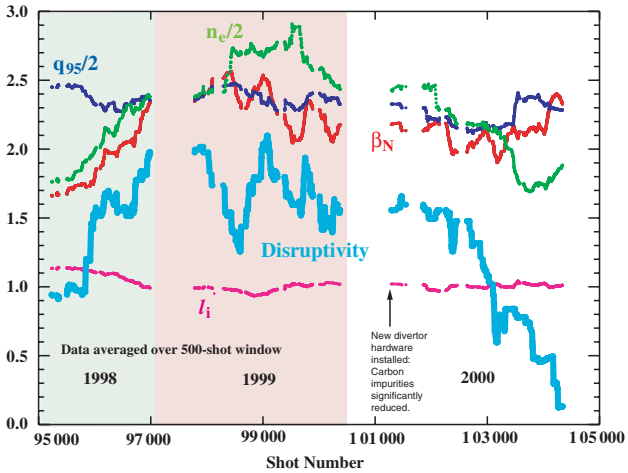

Figure 35. Three-year 'disruptivity' history for DIII-D. Here disruptivity (D) denotes the per-shot likelihood of having a disruption during the current flattop phase of the discharge. $D = 1.0$ corresponds to 10% per-pulse likelihood. There is no evidence of a positive correlation of D with the long-term average values of β_N, l_i or q_{95}. The most evident effect is a long-term decrease in D that followed installation of a new set of divertor tiles. These tiles provided reduced tile edge heating and hence reduced C impurity generation during high-power plasma operation. While there may be evidence of a positive long-term correlation of D with n_e (in units of 10^{19} m^{-3}) the disruptivity data alone does not allow unambiguous identification the underlying mechanism(s) responsible [305].

projections of what may occur in ITER. These projections are subject to caveats that what is seen in present tokamaks vis-à-vis global frequency of disruption (typically about 10% overall, in terms of the disruptions per shot) and the balance between 'physics' and 'external' causes and the breakdowns within these categories are clearly influenced by both the nature of experimental programmes being conducted, the present reliability of hardware systems and the presently limited ability of human operators to fully predict the outcome of experimental campaigns that intentionally push plasma operation close to the one or more of the many operational limits inherent in achieving high-performance plasma operation. In simple terms, this means that campaigns mounted to explore or expand operational boundaries tend to cause more frequent disruption. In the last regard, however, there has been significant recent progress in providing tokamak operators with 'online' impending disruption prediction data that can, in principle, be used to either avoid disruption or to soften the consequences of disruptions that do occur e.g. [306]. The subject of techniques for reliable *a priori* disruption prediction is addressed in section 3.6.

On the converse side, there is also increasing confidence that, given achievement of a well-defined set of plasma operation conditions and reproducible hardware and PFC/wall conditions, there is little evidence for any secular tendency of such plasmas to disrupt as the duration of the steady-state phase is increased. Simply put, there is no evidence for any 'random trigger' mechanism that will arise to cause a disruption in what should otherwise be a sustainable plasma discharge. The 'hybrid H-mode' discharges obtained in DIII-D provide an example of this behaviour; these plasmas exhibit stationary conditions (including stable repetitive sawteeth, Type I ELMs and saturated 3/2 NTMs) for a flattop duration of more than 9 s, or about 9 current profile relaxation times (the relaxation-time-equivalent duration in ITER would be ∼3600 s) [241]; see also chapter 6 of this issue [2] for a fuller discussion of hybrid H-modes. There is no indication that the duration of this type of enhanced-performance discharge is limited by anything other than the pulse-duration capabilities of the tokamak systems needed to produce and sustain it.

3.2. Thermal quench energy loss and deposition

The discussion and data presentation that follows focuses on present understanding and estimates of the energy deposition magnitudes and time scales expected on the ITER divertor and first-wall (FW) PFC surfaces owing to the effects of unmitigated disruptions. Figure 36 illustrates the major plasma magnetic and thermal energy sources and sinks (deposition locations) relevant to consideration of plasma energy flow and deposit during disruption. Estimated magnitudes of the various energies, deposition-affected areas and time scales are given in Table 6 (data from [307]). The purpose of introducing the various considerations and symbolic definitions in figure 36 is to provide a quantitative basis for evaluating present data and making predictions for ITER and beyond.

Figure 36, and the data in table 6, also indicate a number of important considerations that apply to estimates of the resulting ITER divertor PFC and FW energy deposition magnitudes and time scales. These considerations and the present experimental basis for figure 36 and the values in table 6, are discussed below.

Plasma energy magnitudes and thermal energies at disruption. The plasma thermal energy, W_{th} and the in-plasma poloidal magnetic energy W_{mag} for a representative $Q = 10$ full-performance ITER 'ELMy' H-mode (Scenario 2) plasma are estimated to be about 350 MJ and 395 MJ, respectively [307]. The magnetic energy cited in table 6 does not include the additional magnetic energy, external to the separatrix but internal to the torus vacuum vessel that contributes to in-vessel energy deposition during the current decay phase of a disruption (see section 3.3). This energy is estimated to be $\leqslant\sim$200 MJ. While the time scales for the thermal quench and current quench phases of an ITER disruption are predicted to be reasonably distinct (1–10 ms versus $\geqslant$35 ms), there is some potential for overlap and accumulation of deposited energies in the ITER divertor and FW surfaces. Concurrent thermal and magnetic energy deposition on divertor surfaces is frequently observed in present experiments (see below). Furthermore, as table 6 indicates, in ITER, the incremental time-normalized first-wall energy deposition from the current quench phase is not completely negligible (especially if spatial peaking factors and the ex-plasma magnetic energy, not included in the table 6 data, are taken into account) and may be sufficient to bring the surface temperature of the ITER first wall (beryllium) up to or above melting by the end of the disruption. Hence magnetic energy deposition during the current quench phase becomes a significant incremental consideration, especially in high-radiation-fraction thermal quench scenarios.

Pre-disruption thermal energy loss. The IPB thermal deposition guidelines [1] and the estimates in table 6 were (are) based on the conservative assumption that the plasma thermal

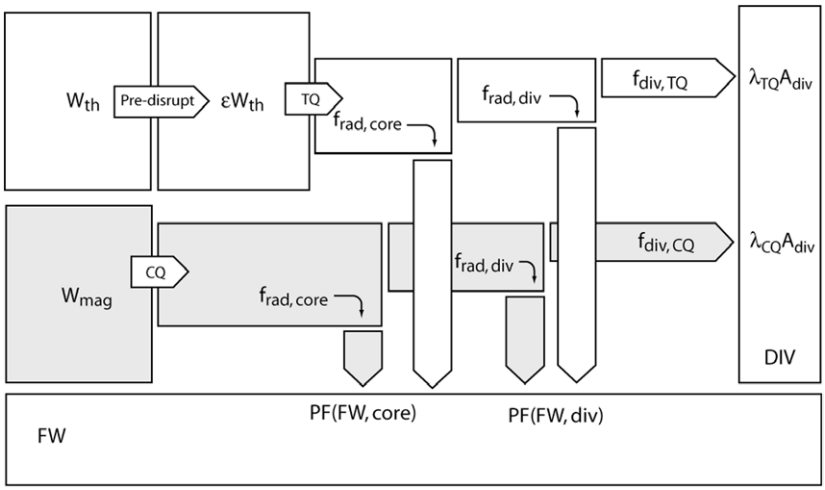

Figure 36. Schematic representation of plasma thermal and magnetic energy flows, and deposition locations during disruption. The possibility of before-disruption (pre-disrupt) loss of W_{th} on an energy confinement time scale (ε factor) is indicated (see text). The PF(FW, core) and PF(FW, div) parameters indicate spatial peaking factors (PFs) relevant to radiative deposition on the overall first-wall surface and/or near the divertor entrance and λ is the broadening parameter of the heat load on the divertor footprint area (A_{div}).

Table 6. Energy-related parameters for ITER disruption (Scenario 2: ELMy H-mode).

Parameter	Symbol	Value	Basis or comment
Plasma magnetic energy	W_{mag}	395 MJ	From equilibrium model; does not include energy external to separatrix
Plasma thermal energy	W_{th}	353 MJ	For $P_{fus} = 400$ MW, with H98$(y, 2) = 1.0$
Plasma surface area	S	683 m^2	At separatrix; FW area (excluding divertor entrance) is similar
Active divertor area (normal operation)	A_{div}	3.5 m^2	FDR physics document
Active divertor area (disruption)	$A_{div,\ dis}$	25 m^2	$\sim 7A_{div}$ (equivalent to $\lambda_{TQ} = \lambda_{CQ} = 7$)
Internal W_{th} relaxation time	τ_2	0.7 ms	*IPB* scaling ($\sim R^1$)
W_{th} loss, W_{div} deposit time	t_{div}	1–10 ms	(1.5–15) τ_2; see text
Current quench time (minimum)	t_{CQ}	35 ms	Based upon $t_{CQ}/S = 1.67$ ms m^2; see section 3.3.1
$U/(A * t^{0.5})$ for divertor	$U*_{div}$	144–446 MJ m^{-2} s$^{-0.5}$	100% W_{th} to divertor; $t_{dep} = t_{div}$ (range is for 1–10 ms)
$U/(A * t^{0.5})$ for FW	$U*_{FW}$	5.2–16.3 MJ m^{-2} s$^{-0.5}$	100% W_{th} to FW; $t_{dep} = t_{div}$; uniform deposition
$U/(A * t^{0.5})$ for FW	$U*_{mag,\ FW}$	3.1 MJ m^{-2} s$^{-0.5}$	100% W_{mag} to FW; $t_{dep} = t_{CQ}$; uniform deposition

energy at the time of disruption is equal to the thermal energy of the 'maximum performance' 'parent' plasma. For ITER, this thermal energy (~350 MJ) would be that of the high-gain ($Q = 10$) plasma needed to realize the ITER fusion power and energy gain (Q) goals. A high-performance steady-state plasma is expected to have a similar thermal energy.

The assumption of disruption at full thermal energy (i.e. $\varepsilon = 1$ in figure 36) is not always borne out by what is seen in present-day 'high-performance' experiments that attempt to obtain maximum plasma thermal energy. For example, in JET, for most 'types' of disruption, the plasma thermal energy at the time of the thermal quench, $W_{th,dis}$, is substantially less than the thermal energy ($W_{th,max}$) of the parent high-performance plasma state that the disruption stems from [308]. With the exception of disruptions in JET arising during high-β_N ITB discharges and of disruptions that occur from the effects of a pure 'hot-plasma' VDE, wherein loss of vertical control allows a full-performance plasma to drift vertically before the disruption occurs, see section 3.3.2, there can be substantial (up to 90%) 'thermal collapse' energy loss before the actual disruption occurs. In contrast, ITB and VDE disruptions typically exhibit $\leqslant$10% before-disruption thermal energy loss.

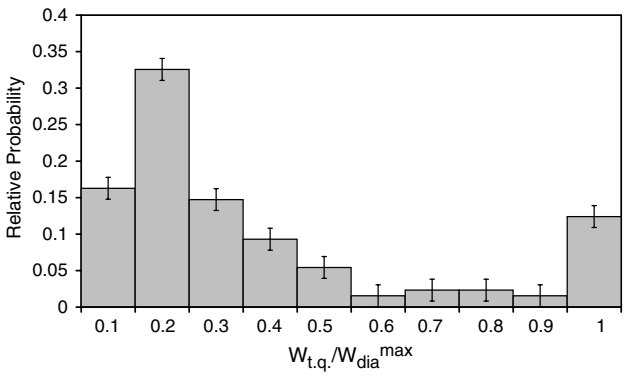

Figure 37. Distribution of pre-disruption thermal energy loss in a 129-example sample of disruptions of JET 'high-energy' plasmas [308]. The ratio of thermal energy at the start of the disruption W_{tq} (denoted $W_{th,dis}$ in the text) to maximum energy in the discharge W_{dia}^{max} ($W_{th,max}$ in text) is plotted [308].

Figure 37 shows the distribution of before-disruption energy loss observed in a representative sample (129 examples) of disruptions occurring in JET 'high-energy' ($W_{th,max} \geqslant 4.5$ MJ) discharges. Approximately 80% of the total have

$W_{th,dis}/W_{th,max} \leqslant 0.5$. The median energy fraction is $W_{th,dis}/W_{th,max} \cong 0.2$. Only 13% of the sample (the ITB and VDE cases noted above) comprise 'full-energy' disruptions with $0.9 \leqslant W_{th,dis}/W_{th,max} \leqslant 1.0$. While the correlation of pre-disruption thermal loss with JET plasma operation procedures has not yet been systematically examined, many of the examples with significant loss are obtained under conditions where 'soft-stop' disruption mitigation measures (decrease of auxiliary heating and elongation and/or plasma current ramp-down) are taken upon initial detection of MHD instability. Given that the delay in many of the 'high-energy' examples between onset of MHD instability or other disturbance and disruption can be ~1 s, these measures are likely a contributing factor to the many low-fraction examples seen in figure 37. 'Natural' MHD instability induced thermal energy loss (confinement degradation) is undoubtedly also an important factor.

Similar magnitudes and relative time scales of pre-disruption W_{th} loss are seen in DIII-D and ASDEX Upgrade, often owing to the effect of slowly growing or saturated NTMs, or, in DIII-D, owing to the effect of slowly growing (marginally stabilized) RWMs. Deterioration or loss of H-mode confinement following approach to the Greenwald density 'limit' can also produce a similar 'slow' confinement-loss thermal collapse that precedes disruption. This type of 'density limit' disruption, which typically results from excessive plasma core and edge radiation, is often described as a 'cold-edge' or 'radiation limit' disruption. The relative frequency of the occurrence of cold-edge disruptions preceded by a H-mode to L-mode transition, and a subsequent slow thermal collapse, remains to be quantified statistically.

The JET and other tokamak experiences clearly demonstrate that substantial plasma thermal loss—on a time scale commensurate with the energy confinement time of the parent plasma—can occur before disruptions that are initiated by slowly growing MHD instabilities. The counter examples in JET occur in plasmas where fast-growing MHD instabilities cause a 'prompt' disruption. Similar 'prompt disruption' examples are well known and reported in DIII-D and ASDEX Upgrade. How these 'fast' versus 'slow' energy loss behaviours and operational experiences will extrapolate to ITER (wherein plasma control capabilities and control time scales may differ from those now available in present experiments) remains to be determined. But there does appear to be a prospect for using the natural properties of slowly growing MHD-initiated thermal collapse, perhaps supplemented by active 'soft-stop' procedures analogous to those employed in JET, to reduce ITER plasma thermal energy levels prior to otherwise unavoidable disruptions. The potential benefit of this type of reduction in ITER is considered below.

Plasma operation in JET with a reversed shear configuration seems so far to be the exception to the experience with slowly growing instabilities: as the JET experience shows, and is well-known in other tokamaks, disruptions stemming from such advanced scenarios are often prompt, with very little pre-disruption energy loss. Thermal and current quench rates in reverse shear cases are also among the fastest examples seen in a given device (see discussion of thermal quench times below and current quench time discussion in section 3.3.1). Examples of similarly prompt disruptions are also seen in DIII-D in positive-shear H-mode plasmas wherein NTM suppression or avoidance allows operation near or above the ideal 'no-wall' MHD stability (Troyon) beta-limit.

Duration of the thermal quench. The thermal quench is a phenomenon in which most of the plasma thermal energy is lost by the plasma and is subsequently deposited, by conduction or/and convection and radiation, on the limiter and divertor surfaces. The duration depends on the machine size; it is of the order of tens of microseconds in small tokamaks, hundreds of microseconds in medium-sized tokamaks and has been observed to reach a few ms in JET. The duration is also observed to vary significantly within a given machine, apparently depending on the 'type' of disruption and/or the nature and growth rate of the triggering MHD instability. There is also evidence that the pre-disruption magnetic shear profile (positive $dq/d\psi$ versus weak or strongly reversed $dq/d\psi$) can affect the rapidity of the thermal quench duration.

The PFC energy deposition phase of the thermal quench is typically preceded by one or more internal energy redistribution phases, wherein a rapid redistribution of the energy within the plasma occurs. Cases exhibiting one, two or three or more sequential redistributions are known. References [1] and [309] describe in detail a typical two-stage thermal quench. The thermal energy is firstly redistributed within the plasma inside the $q = 2$ surface; after a delay $\tau_{(1-2)}$ the energy barrier within the closed flux surface region breaks and the energy is redistributed between the plasma and the open flux surface region within a time τ_2. It should be noted that in some cases these phases effectively merge and during the initial redistribution of plasma energy there can be a significant power outflux from the confined plasma [308,310]. In the time interval τ_2, the energy starts being lost to the divertor. Data from the rudimentary multi-machine IPB disruption database (see chapter 3 of [1]) suggest that τ_2 scales with the minor radius of the device and that for ITER, $\tau_2 = O(1)$ ms. More recent data from ASDEX Upgrade and JET do not contradict this estimate for the relaxation or loss time.

The observed relaxation and loss times are qualitatively consistent with the expected effects of global ergodization of the internal plasma magnetic surfaces and/or the development of large-scale internal ballooning mode structures (see section 3.1, above). In all cases, the internal relaxation and loss times are much faster than the predictions for the time scale of resistive reconnection. The theoretical basis for the time scale of the fast quench and internal redistribution(s) are not well understood, and a validated basis to predict the ensuing rate of plasma energy loss and particle transport to the PFC surfaces is not yet available.

Time scale of PFC energy deposition. Specifying a PFC deposition time equal to the plasma loss time gives an upper bound to the possible time-normalized PFC energy deposit. Examples of such 'loss-rate-determined' energy deposit have been observed [1]. However, ASDEX Upgrade data show energy deposition times $> \tau_2$, as illustrated by high-time-resolution infra-red (IR) camera data for the divertor target power loading during a disruption (figure 38). These data show that the rising phase of the heat pulse on the PFCs,

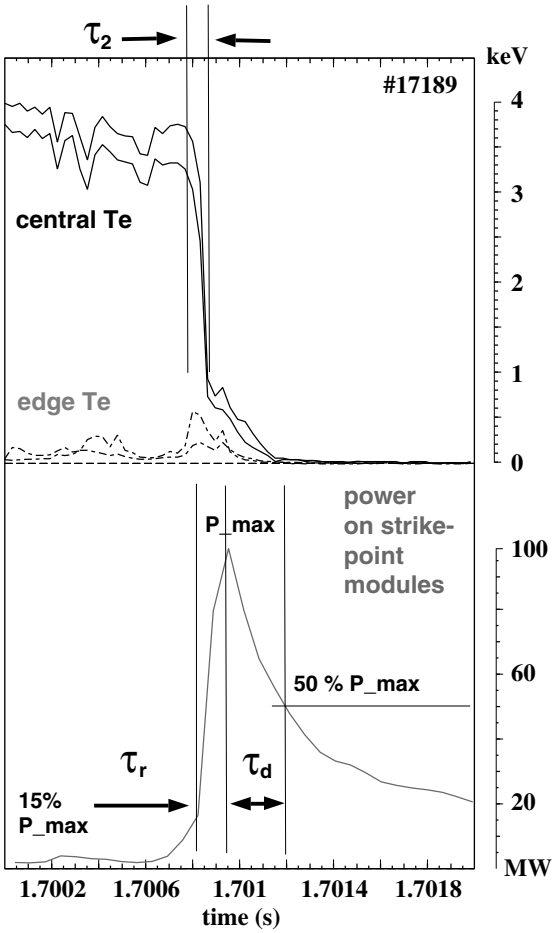

Figure 38. Plasma temperature and divertor target power data from ASDEX Upgrade.

τ_r, coincides approximately with the fast-drop phase of the thermal quench, that is, with τ_2. The actual amount of energy that reaches the target in the fast-rise phase is small: most of the energy arrives at the target plate in a subsequent phase and on a longer time scale, τ_d. The overall energy deposition versus time waveform is also 'triangular' rather than 'square': there are time-waveform shape factors of order unity applicable to evaluating the magnitude of the resulting time-normalized surface heating parameter.

The waveforms in figure 38 give some indication of the mechanisms involved. The initial central T_e loss event of duration τ_2 (here resulting in a fast central energy drop of $\sim$80%) gives rise to a corresponding dynamic increase in edge T_e, which in turn is apparently responsible for the rising phase of the divertor target power and the initial part of the subsequent power decay. The central T_e decay has a 20–0% 'tail' that follows the edge temperature decay. The data strongly suggests that the initial global mixing (τ_2) event is followed by a more-or-less uniform decay of plasma energy to the divertor on a time scale that is (in this example) $\gtrsim\sim 5\tau_2$. There is also an even-longer power decay 'tail' that continues after the measured edge and central T_e are essentially zero.

For ASDEX Upgrade, the measured τ_r is sometimes as short as the time resolution of the IR system, i.e. 0.1 ms, and is, on the average, 0.5 ms; τ_d is longer than the time resolution of the IR system and is on the average, about 1 ms [311]. This total heat-pulse duration is significantly longer than the IPB scaling (chapter 3, figure 54 of [1]) prediction of $\tau_2 = 0.2$ ms for ASDEX Upgrade. The ECE-measured central T_e loss time is comparable to the IPB scaling for τ_2.

The physics determining the time scale of the heat pulse duration can likely be described by (i) conduction parallel and perpendicular to the open magnetic surfaces of the SOL, which is affected by SOL turbulence and stochasticity and (ii) the limitations imposed by the electron power flux through the sheath at the plasma–wall interface. These physics elements will certainly contribute to producing the divertor heat pulse 'time-stretching' (integrating) effect that is clearly evident in the ASDEX Upgrade data. The strong non-linearities inherent in a coupled SOL conduction + sheath limit model might also explain the seemingly wide degree of variance seen in plasma energy loss time and heat pulse duration data in a given machine. The role of high-plasma edge density in effecting both 'thermal isolation' (via globally reduced SOL conduction/convection and a reduced divertor target sheath flux limit) and pre-divertor radiative dissipation (to the divertor baffles and/or divertor entrance first wall) in the SOL, of the plasma core thermal energy, is also likely significant. These expectations about thermal energy transport and deposition have been compared with B2 SOLPS code calculations for ASDEX Upgrade [312]. Systematic study of the correlation of plasma loss and target energy deposition times on a multi-machine or database basis remains as a future R&D task that will likely require improvement in W_{th}(r, t) and FW surface energy deposition diagnostics.

Finally, there is also the matter of the actual magnitude of τ_2. In this regard, instrumental limitations in at least some thermal loss data cited in the IPB (mainly based upon SXR emission) may have prevented accurate observation of the fastest thermal loss times. For example, newly available fast ECE measurements [308] in JET of the central plasma thermal energy loss during 'prompt' ITB disruptions show that τ_2 in such disruptions can be as short as 50 μs (cf $\tau_2 = 200$–500 ms reported—based on SXR data—for 'typical' JET plasmas in the IPB).

Extrapolation by the minor radius scale factor of these JET 'ITB-type' thermal quench times to ITER yields $\tau_2 = 100\,\mu$s (cf 700 μs for the standard IPB scaling extrapolation). But the 'time-stretching' factor inherent in setting the time scale of the actual divertor target heat pulse may obviate most of the potentially damaging effect of such very-fast plasma thermal energy redistribution events. Better understanding of all of the physics considerations involved in setting the time scale of transport of thermal energy to the PFC surfaces is clearly needed.

Surface area and locations of heat deposition. Time-resolved power deposition during a discharge is typically measured with IR cameras. Several divertor tokamaks (DIII-D, ASDEX Upgrade, JT-60U, JET) have cameras positioned to view the primary or/and secondary divertor targets at one or two toroidal azimuths, and the heat flux in the divertor has been diagnosed in a variety of discharges. Routine IR monitoring of the first-wall (other than the portions in direct view of the divertor camera) is less common though there are recent developments in the IR measurements of an appreciable fraction the whole

Table 7. SOL power and energy deposition widths in ASDEX Upgrade and JET during the thermal disruption (Λ_{TQ}) and normal operation (Λ_{SOL}); widths are quoted in the outboard mid-plane. The major radii quoted are the geometric radius.

	R (m)	Λ_{max} (cm)	Λ_{TQ} (cm)	Λ_{SOL} (cm)
ASDEX Upgrade	1.65	5	>5	1
JET	2.96	5	$\gg 5$	0.5

plasma-facing surface in MAST [313], JET [314] and ASDEX Upgrade [315], that show evidence of disruption power fluxes outside the main divertor region. TEXTOR [316, 317] is equipped with an IR camera monitoring the limiter.

The evolution of the temperature and density profiles in the plasma SOL or, later on in the disruption, in the halo region, is usually not diagnosed directly (e.g. with Langmuir probes or Thomson scattering) during and after the thermal quench. Nevertheless the IR-measured divertor power deposition profiles (magnetically extrapolated back to the outside midplane SOL) indicate that during the thermal quench, the SOL expands considerably, relative to the similarly observed before-disruption 'steady-state' SOL width, Λ_{SOL}. Observation of SOL expansion factors (λ_{TQ}, see figure 36) of 1–10 in limiter and divertor tokamaks is reported in the IPB [1]; more recent data from medium- and large-scale divertor tokamaks (see below) indicates that the range of SOL expansion factors is shifted to higher values: e.g. $5 \leqslant \lambda_{TQ} \leqslant \sim 20$. With the exception of TEXTOR, the width of the power channel in the SOL during thermal quench is typically larger than the field-of-view width, Λ_{max}, monitored (viewed) by the IR camera (table 7).

In the case of JET, the energy balance calculation (amount of energy deposited within the IR camera field-of-view) indicates that most of the plasma thermal energy is transported beyond the 5 cm region of the SOL 'seen' by the IR camera and new IR data indicates that at least some of this energy is conducted to the walls [314].

Magnetic energy deposition. The phenomenology and physics basis for the current decay phase of a disruption (wherein the in-vessel plasma magnetic energy is dissipated) are treated in detail in section 3.3. Here only the magnetic energy deposition aspects will be considered. While recent (since the IPB) data still generally support the recommendations given in the IPB about magnetic energy deposition (80–100% of W_{mag} more-or less uniformly radiated to the FW, 0–20 % locally conducted to the divertor or divertor baffles) [318], experience in present tokamaks shows that a significant fraction of W_{mag} can be (is, in at least some cases) deposited on the divertor targets during or following completion of the thermal energy quench phase. Conversely, given the present observations that a significant fraction of W_{th} is being deposited somewhere other than in the IR-observed portions of the divertor, deposition of the 'missing' fraction of W_{th} on the divertor entrance baffle surfaces and/or at least a limited portion of the FW must be taking place.

Despite the fact that a large fraction of the plasma magnetic energy is radiated during current quench, temporally and spatially localized power deposition has to be expected in a ITER current quench, as already observed in the existing tokamaks. In JT-60U, the power deposited during the runaway current termination was observed to be spatially localized (within a poloidal width of 20 cm on the inner divertor plate) and to occur in pulses of the duration of less than 0.25 ms [319] (see also section 3.4). In both ASDEX Upgrade [320] and JT-60U [321], the occurrence of the $n = 1$ asymmetry of the halo currents is accompanied by a MHD phenomenon expelling some particles and residual thermal energy from the plasma in a time interval comparable to the thermal quench time.

During the current quench, part of the plasma magnetic energy is transferred to the conducting coils and structures around the plasma. The physics of mutual induction and dissipation of the current in conducting structures around the machine is known. Recent evaluations of the amount of energy dissipated in the external conductors in JET [322] show that typically 35% of the magnetic plasma energy is dissipated in conductors outside of the plasma during disruptions. In ITER, owing to the low-resistance and long ($\sim$0.7 s) $n = 0$ time constant of the ITER vacuum vessel, transfer of in-vessel magnetic energies to external conductors is expected to be negligible in most fast disruptions and fast VDEs. However, slow disruptions and/or slow VDEs may have some in-vessel energy transfer.

Energy balance. In JET the thermal energy deposited in the divertor within the time scale of the thermal quench (1 ms) is only a few % of ΔE_{th} [323] and is spread on the whole divertor. The lack of full accountability for the plasma thermal energy during the thermal quench cannot be fully explained by (1) the effect of radiative losses upstream of the divertor, (2) asymmetries in the on-divertor distribution and (3) uncertainties in the evaluation of the divertor surface power fluxes. In the case of JET a large perpendicular transport of energy during the thermal quench and losses beyond the field-of-view of the IR camera were postulated to justify the experimental observations; more recent observations support this postulate [314].

The amount of energy deposited on the divertor during the whole disruption in JET amounts to up to 10% of the total pre-disruption energy, $\Delta E_{tot} = \Delta E_{th} + \Delta E_{mag}$. Between 50% and 80% of the total energy is found to be radiated. This fraction does not seem to depend on the type of disruption. ITB collapses can generate the highest and most localized power densities.

In DIII-D [324] the energy conducted to the divertor within the whole disruption is 15–50% of the total plasma energy or 50–100% of the thermal energy of the pre-disruptive plasma, apparently less in the case of radiative-limit disruptions and more in the VDE and beta-limit disruptions. Similar results are reported in [310] for ASDEX Upgrade. During the 4 ms centred about the thermal quench time, the energy deposited on the lower divertor is in average 90% (and can reach 200%) of the thermal energy. This suggests that during this time already a fraction of the magnetic energy may be dissipated in some cases. The amount of energy deposited on the lower divertor during the whole disruption is in average 30% (and can reach 45%) of ΔE_{tot}.

Discussion, conclusions and recommendations for future R&D. A schematic and parametric summary of the present

Table 8. Single-parameter mitigation factors to avoid carbon ablation.

Mitigation method	Pre-disruption W_{th} reduction	Pre-divertor radiation	Time stretching	SOL broadening
Symbol (figure 36)	ε	1-$f_{div,\ TQ}$ ($f_{div,\ core}$ + f_{div})	$\gamma * \tau_2$	λ
Reference basis	1.0	0.0	$\gamma = 1.5$	$\lambda = 7$
Requirement	$\leqslant$0.09	$\geqslant$0.91	$\gamma \geqslant 86$	$\lambda \geqslant 80$

status and expectations (predictions for ITER) are embodied in figure 36 and table 6. The most immediate change since the IPB lies in the 2-fold reduction in the ITER area-normalized and time-normalized plasma thermal and magnetic energy loadings. With the corresponding dimensional scalings of the divertor deposition areas and plasma energy loss (τ_2) time scales taken into account, the time-normalized W_{th} loading on the divertor targets drops from $\sim$1000 MJ m^{-2} s$^{-0.5}$ for the ITER-EDA design to $\sim$450 MJ m^{-2} s$^{-0.5}$ for ITER. This latter loading, which is calculated for a full-energy disruption ($\varepsilon = 1$) with 100% W_{th} into the divertor, a 7-fold expansion of the deposition width ($\lambda_{TQ} = 7$) and with the further conservative assumption of prompt thermal deposition ($t_{dep} =\sim 1.5\tau_2$, with τ_2 following the IPB scaling), can be regarded as an upper bound to the expected ITER divertor energy loading. While this 'upper bound' estimate is still a factor of $\sim$12 ($\sim$8) above the ablation (melting) onset threshold for carbon (tungsten), the prospects for obtaining either reduced or negligible divertor target erosion are now enhanced relative to the same basis situation for the EDA ITER. The authors of [308, 325] reach similar conclusions about the range of possible divertor energy loadings in ITER; it should be noted in [325] that a 75% reduction from the full ITER $Q = 10$ scenario plasma energy is assumed before the thermal quench (i.e. $\varepsilon = 0.25$). Table 8 demonstrates in a parametric manner the 'single-parameter' mitigation factors required to reduce the $\sim$450 MJ m^{-2} s$^{-0.5}$ loading to the $\sim$40 MJ m^{-2} s$^{-0.5}$ threshold for carbon ablation.

The single-parameter mitigation requirements embodied in table 8 can be compared with the corresponding observations in present tokamaks. Pre-disruption energy loss factors as large as 0.1 have been observed in JET. Very large scrape-off layer (SOL) broadenings, corresponding to $\lambda > \sim$20 and/or very high 'pre-divertor' radiation fractions (>90%) can be inferred from the low-divertor W_{th} energy depositions observed in JET. Divertor heat pulse duration 'stretching' by a γ factor of 5–10 is routinely observed in ASDEX Upgrade.

Except for the JET observations of high fractions of pre-disruption W_{th} loss in 'slow' density limit and similar disruptions (and equivalent observation in ASDEX Upgrade), none of the single-parameter mitigation observations explicitly yet matches the ITER 'no-ablation' requirement. On the other hand, a modest combination of pre-divertor radiation ($f_{div,\ core}$ + f_{div} = 0.4) moderate time-stretching (γ = 10) and doubling the SOL broadening (λ = 15) would reduce the ITER divertor target deposition to $\sim$48 MJ m^{-2} s$^{-0.5}$. This comes close to meeting the no-ablation or no-melting thresholds (but our estimates here should not be interpreted as constituting actual assessments of ITER divertor target response). Before-disruption energy loss would further reduce the time-normalized loading. While the present understanding and data basis for W_{th} deposition and mitigation approaches is insufficient to make definitive conclusions about what will happen or will be possible in ITER, present data and the considerations developed above suggest that nearly melt- or vaporization-free divertor PFC operation may be obtainable in at least some of the proposed ITER plasma operation regimes, especially in plasma operating regimes or modes where natural or 'soft-stop' before-disruption thermal energy loss will be possible.

We also note here that while there has been modest progress since the compilation of the ITER Physics Basis, there remains significant work to do in terms of our ability to account for the magnitude and time history and spatial distribution of the deposited thermal and magnetic energy, model and simulation development and testing of mitigation techniques. Given the importance to ITER (and DEMO) of minimizing, or eliminating, the operation-limiting effects of disruptions this work remains a key area for future studies. ITER must prove high-reliability disruption avoidance and mitigation methods, essential for DEMO (where disruptions power loadings problems will be more severe than for ITER).

3.3. Current quench dynamics

During the current quench phase of a disruption in a vertically elongated plasma, the decay of the plasma current and the ensuing motion of the plasma column induce toroidal currents in the nearby toroidally conducting structures (e.g. the ITER torus vacuum vessel) and also drive force-free helical current flow in the wall-contacting ex-plasma 'halo' region that lies beyond the last closed plasma flux surface. Both the induced toroidal current and the ex-plasma halo current flow act to mediate the dynamic evolution of the current channel during the current decay phase. The components of the ex-plasma halo current that reconnect through the vessel and the in-vessel conducting structures (e.g. the ITER shield blanket modules the divertor baffle and cassette modules) give rise to in-vessel halo currents that produce forces on these structures. In addition, the plasma current decay and motion induce locally circulating 'eddy currents' in nearby plasma-facing conducting structures (e.g. the ITER first-wall and shield blanket modules), and the interaction of these induced circulating currents with the toroidal and poloidal fields gives rise to localized torques and overturning forces on these PFC structures.

The magnitude and time-history of the eddy current loading depends on the local rate of change of the poloidal magnetic field, B_p, and on the eddy current decay time constant (L/R time) of the affected structures. Hence understanding of the range of possible plasma current quench rates, local dB_p/dt rates and the simultaneous magnitude and direction of halo current flow in the affected structures is needed. Simple extrapolation of the quench time based upon a linear current decay waveform and the assumption of a static (fixed in space) plasma current channel may not necessarily provide

the accurate assessment of dB_p/dt needed, for example, for evaluation of the loadings on the ITER in-vessel shield modules. Hence there has been recent attention (since the writing of the IPB) to the details of the time history of the plasma current decay waveform and also to the eddy current and halo current generating effects of motion of the plasma current channel during the current decay. In section 3.3.1 that follows, emphasis is on global current decay; effects of current channel motion due to a VDE, are considered separately in section 3.3.2. Ultimately, the combined effects of current decay and motion and the generation of poloidally flowing halo currents in in-vessel structures must be assessed in a self-consistent manner using an integrated dynamic model (see section 3.5).

3.3.1. Global toroidal current decay. Evaluations detailed in the ITER Physics Basis, chapter 3.4 [1], of the range of plasma current decay rates observed in seven circular and elongated cross-section tokamaks (Alcator C-MOD, ASDEX Upgrade, DIII-D, JET, JT-60U, TFTR and Tore Supra) resulted in the finding that the lower bound on the inferred area-normalized current decay time, $\Delta t_{60} = t(80\%)–t(20\%)$, normalized by the elliptic-approximation estimate for the plasma cross-section area, $S* = \pi\kappa a^2$, was $\Delta t_{60}/S* = 0.8\,\mathrm{ms\,m^{-2}}$, is independent of plasma elongation and average toroidal current density. If this area-normalized time for the inferred 60% current decay (from 80% I_{p0} to 20% I_{p0}) is directly extrapolated to the time for 100% current decay, the resulting 'linear-basis' minimum decay time is $0.8 * 100/60 = 1.33\,\mathrm{ms\,m^{-2}}$. It should be noted that the basis for determining Δt_{60} in the IPB was not uniform between tokamaks and for example the JT-60U value is inferred from observed maximum I_p quench rate.

The shape of the current waveform was not explicitly considered in chapter 3.4 of [1], although anecdotal and indirect evidence for many discharges having an 'S-shaped' decay waveform with a nearly linear intermediate current decay phase was cited. More recent systematic analysis of JT-60U current quench waveforms shows that they can be fitted with either an exponential-like or a linear-like function [326]. Exponential fitting is better when the area-normalized quench time is small; linear fitting is better otherwise. Differences between linear and exponential models for the current decay waveform (and hence the corresponding estimates of dB_p/dt) affect eddy current forces in structures with time constants comparable to or shorter than the current decay time. Calculations of global magnetic pressures or eddy current forces in long-time constant structures are not affected by dB_p/dt, since such pressures or eddy current forces depend on the total change in poloidal field, independent of time scales.

Figure 39 demonstrates the nearly exponential nature of the central (80–20%) segment of a typical fast JT-60U current decay waveform and illustrates how approximating the decay by a linear extrapolation, determined by taking the time difference between 80% current and 20% current, results in an approximately 2-times underestimate of the peak rate of current decay. Elementary considerations show that the underestimating factor for an exact exponential waveform is 2.3; the actual JT-60U waveform yields an ~40/18 (= 2.2) ratio of linear versus exponential decay times.

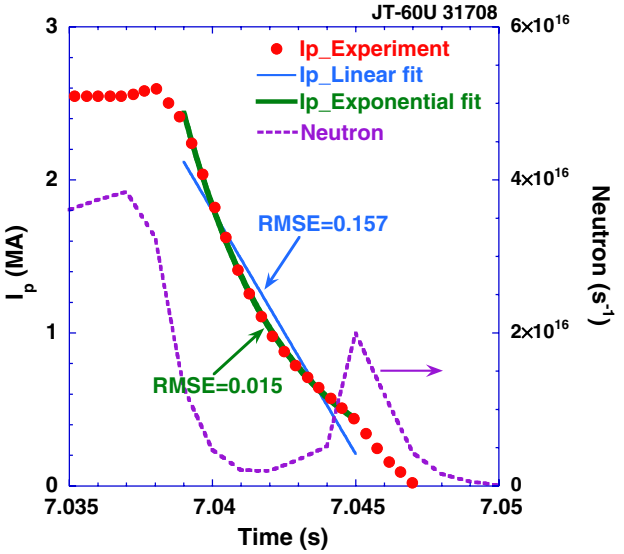

Figure 39. Waveforms for a typical fast current quench observed in JT-60U. Deviation of the final portion of the decay from a pure exponential is attributed to loss of runaway current. The corresponding increase in neutron rate is attributed to photo-neutron production [327].

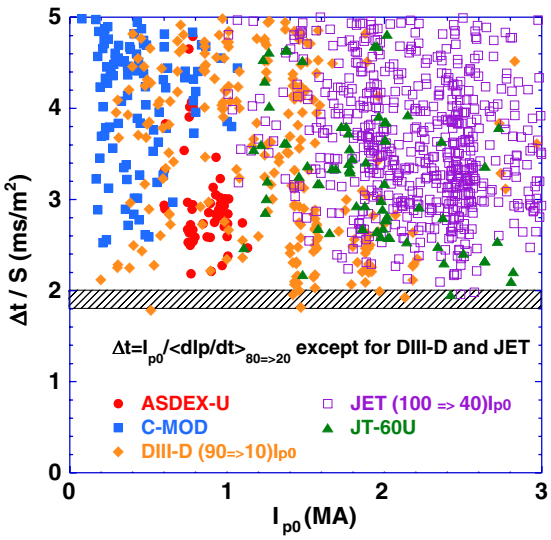

Figure 40. IPB data, replotted in terms of the linearly extrapolated current decay time, Δt, normalized by their respective poloidal cross-section area, S [327].

Figure 40 shows, for non-circular tokamaks only, current quench data from the original 1996 international database (IPB chapter 3, section 4 of [1]), supplemented with updated JT-60U data [327] based on the actual 80–20% quench time for I_p. In figure 40, the data are plotted to show the linear extrapolation quench time, $dI_p/\langle dI_p/dt\rangle$, calculated, except as noted for JET and DIII-D, on an 80–20% of I_{p0} basis. Data for DIII-D and JET are included, corrected for the respective data basis; 90–10% decay for DIII-D and 100–40% decay [308] for JET.

For this revised data set, the lower bound on area-normalized linear quench time for 100–0% decay is $1.8\,\mathrm{ms\,m^{-2}}$; this lower bound is about 35% greater than the $\sim 1.33\,\mathrm{ms\,m^2}$ lower bound (100–0% basis) on area-normalized linear quench time inferred from the original IPB-era data

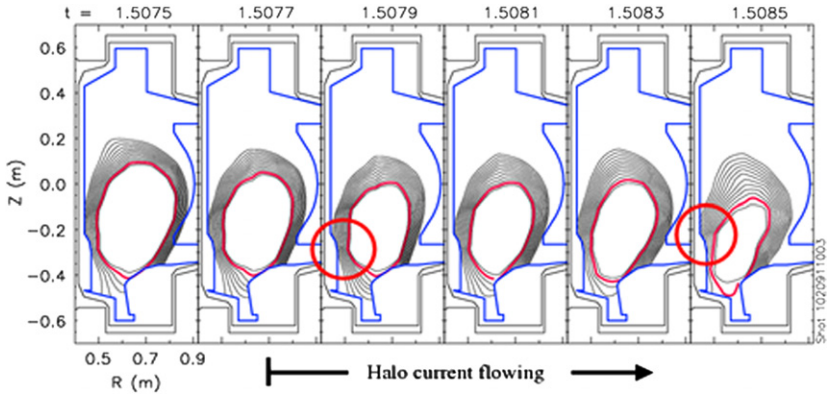

Figure 41. Plasma configuration evolution and halo current region development in a typical Alcator C-Mod disruption. The data shows that much of the ex-plasma halo current region (inferred from magnetic reconstruction data) now misses the new, less-protruding inboard divertor (see discussion below).

(chapter 3.4 of [1]). This increase in quench time is largely due to the JT-60U data being based on an 80–20% I_p quench time (as opposed to the maximum I_p quench rate as in the IPB). Application of the 1.8 ms m^{-2} lower bound to the present ITER, with $S* = 22$ m^2 yields a 40 ms linear quench time, or equivalently, an exponential decay with an 18 ms time constant.

The DIII-D data in figure 40 are based on plasma current decay data from an external (outside the vacuum vessel) Rogowski loop. Recent consideration [328] of decay rates evaluated using an in-vessel 'pseudo-Rogowski' diagnostic (synthesized from a set of in-vessel magnetic probes) indicates that the fastest DIII-D current decay data in figure 40 have $S*$-normalized decay times that fall in the range 1.4–1.7 ms m^{-2}. Further consideration of these DIII-D results and new current quench data, to be collected on a systematic multi-machine basis, under the aegis of the International Tokamak Physics Activity, remains as an ongoing current quench research issue with the most recent analysis showing the lower bound on area normalized quench time should be taken as 1.7 ms m^{-2} [329].

Calculations of the electromagnetic effects of induced eddy current loads in the ITER shield blanket modules are described in [327]. The effects of both a 40 ms linear current decay waveform and an 18 ms exponential decay waveform are considered. The effects of the projected in-vessel halo currents are also included. The conclusion reached therein is that the structural design of the blanket modules is sufficient to withstand the maximum collective loading, albeit with only a small (∼20%) margin relative to structural design allowables. The authors of [327] note that a somewhat slower current decay scenario (one with a 60–80 ms linear decay) would, with the corresponding decrease in eddy current loading, yield a significant increase in margin relative to the structural allowables.

3.3.2. Vertical instability, halo currents and mechanical forces. Since plasmas in typical elongated cross-section tokamaks are inherently unstable against vertical displacements, a sufficiently large and fast change in plasma parameters can cause the loss of the vertical position control, leading to an uncontrolled upward or downward excursion of the plasma column. A true vertical displacement event (VDE) begins with a loss of vertical stability that develops before any appreciable cooling of the plasma centre occurs. Such events are typically described as a 'hot-plasma' VDE. The plasma current centroid moves vertically away from its equilibrium position and the moving plasma column eventually contacts a limiting surface. The direction of initial movement, even in single-null plasmas, can be either towards or away from the X-point, depending on the changes in the plasma current and pressure profiles (i.e. l_i and β_p) as well as the initial location of the plasma current centroid [330].

In a VDE, the plasma continues to move into the wall, reducing the plasma area, typically with little change in the total plasma current, thus reducing the edge safety factor. A pronounced ex-plasma halo current flow also develops (see the discussion and figure 41 below for a representative example of plasma configuration evolution and halo region development during a vertically unstable disruption in Alcator C-Mod), and the 'halo' currents flowing in the wall-contacting ex-plasma region reconnect through the structures that the plasma comes in contact with. The resulting in-vessel currents, which typically flow mostly in a poloidal direction, are commonly called 'halo currents'. With suitable instrumentation, these in-vessel halo currents can be measured directly and the torus-circumference sum of the measured or inferred poloidal in-vessel current flow is typically taken as the measure of the total (in-vessel) halo current, denoted symbolically in the IPB and herein as I_h. Unless otherwise noted, reference to the *poloidal* component of the halo current, in the in-vessel conducting structures, should be understood.

When the plasma is in contact with the wall in-vessel halo currents are also driven to compensate for the plasma toroidal flux lost because of the plasma area reduction. The plasma current stays approximately fixed during the hot plasma–wall contact phase because the resistive decay time of the still hot core plasma is long compared with the vertical motion timescale. When the boundary safety factor decreases to a sufficiently low value (typically less than 2), rapid growth of MHD activity ($n = 1$) produces a fast thermal quench similar to those observed in major disruptions. As the plasma current starts to decay, toroidal currents are induced in the halo region to conserve the poloidal flux [331].

Vertical instability can also occur following the onset of a disruptive thermal quench. In this case, the resulting vertical evolution can be described as a vertical disruption, or more

correctly, as a vertically unstable-disruption (VUD) current quench. The resulting 'cold-plasma' current decay phase shares the most of the attributes—including the generation of in-vessel halo currents and large transient vertical forces on the vacuum vessel—of a 'hot-plasma' VDE. Figure 41 shows magnetic reconstruction data from a typical VUD example in Alcator C-Mod. Similar examples of plasma evolution and halo development during disruption or VDEs are obtained in ASDEX Upgrade, DIII-D, JET, JT-60U, etc.

It is noted here parenthetically that both the IPB-era evaluation of the lower bound on area-normalized current decay rates (chapter 3.4 of [1]), and the more recent reconsideration of IPB and post-IPB data for elongated tokamaks [327] ignore any distinction between decay rates (following onset of thermal quench) for hot-plasma VDEs and vertically unstable cold-plasma VUDs. In the IPB comparison of lower-bound data from circular (vertically stable) and elongated (vertically unstable) tokamaks, there was no observable distinction in the lower bound on minimum area-normalized current quench times.

Halo currents: phenomenology and physics basis The basis and phenomenology of the generation of both ex-plasma and in-vessel 'halo currents' is described in the IPB (chapter 3.4 of [1]) and in, for example, [330, 332]. Briefly put, from a tokamak design basis point of view, the most important considerations for halo currents are (1) the magnitude of the in-vessel halo currents, (2) the toroidal asymmetry of these in-vessel halo currents and (3) the resulting $\boldsymbol{j}_{\text{in-vessel}} \times \boldsymbol{B}$ forces that act on the vessel components. It is the sum of these local forces (plus the global vertical forces that are directly generated by induced toroidal vessel currents) that give rise to the large vertical vacuum-vessel forces that are observed in present tokamaks during VDEs and VUDs. Toroidal asymmetry in the halo current distribution can also give rise to an asymmetric radial de-centering and/or tilting forces on vacuum vessel systems. Finally, the magnitude and spatial (radial) extent of the ex-plasma halo current flow becomes the dominant factor (especially in the end phase of the current decay) in determining the equilibrium dynamics of both hot-plasma VDEs and cold-plasma VUDs. More simply put, in this end phase, most of the remaining plasma current flows in the halo region. The presence and dynamics of this ex-plasma halo current must be self-consistently taken into account in calculations of the plasma equilibrium evolution and the resulting vacuum vessel and in-vessel-component toroidal eddy currents, and forces on these component systems. The need for—and development of—self-consistent models for plasma dynamics and the resulting ex-plasma and in-vessel halo currents are further addressed in section 3.5.

Recent systematic studies of halo current characteristics and phenomenology have contributed to developing the improved physics basis understanding needed for such models. Since the magnitude of the poloidal halo current flow depends explicitly on the plasma edge (last closed flux surface) safety factor, q_ψ, a database for the edge safety factor at the onset of thermal quench during vertical displacement events has been compiled in JT-60U [332]. It is found that the edge safety factor at thermal quench onset in hot-plasma VDEs varies between 1.5 and 2, depending on the discharge conditions. This finding confirms the less-explicit understanding documented for the *IPB* that thermal quench onset in VDEs typically occurs for $q_\psi \sim 1.5$–2. The ability of the plasma core to maintain gross MHD stability at $q_\psi < 2$ is also consistent with the expected stabilizing effect of the current-carrying ex-plasma halo. In JET, higher halo currents and asymmetries are observed when the boundary safety factor decreases close to unity [333]. However, events starting with a 'thermal collapse' usually have much smaller halo current fraction and hardly any asymmetry. This is believed to be because in these cases, the plasma current decay is faster than the plasma cross section reduction and a relatively high-boundary safety factor is maintained.

The vertical force due to the plasma vertical displacement is proportional, to first approximation, to the product of the plasma current, the vertical displacement of the current centroid and gradient of the before-disruption external equilibrium field. Therefore a substantial reduction in vertical load on the vessel can be achieved if the vertical displacement of the plasma during current quench can be minimized. The concept of effecting plasma operation in a weak single-null or balanced double-null plasma with a before-disruption current centroid positioned near the vertical-instability 'neutral point' has been advanced [334] as being advantageous to minimized vertical loads. However, the electromechanical loads due to the plasma current decay would be still present, as well as the risk of generating runaway electrons. In addition, concerns have been expressed that variations in the before-disruption or during-disruption evolution of the vertical-stability-determining parameters (l_i and β_p) may make it impossible to determine, for a single-null plasma, a single neutral point location (height) that avoids vertical instability for a range of disruption 'types' [330].

Halo current magnitude, asymmetries and experimental observations. Incorporation of new halo current measurements into the existing *ad hoc* ITER disruption database has been ongoing since the ITER Physics Basis was published. These new data additions reflect the effect of several major hardware modifications, including, in particular, ASDEX Upgrade operation with the Divertor II (Lyra) and Divertor IIb target and baffling modifications, and Alcator C-Mod operation with a newly shaped inboard divertor. The new data also include new halo current measurements from JET that reflect refurbishment of disruption-related instrumentation, and halo current measurements from the low-aspect-ratio MAST spherical tokamak. The resulting updated plot of toroidal peaking factor (TPF) versus $I_{\text{halo}}/I_{\text{p0}}$, (figure 42), includes new data from JET [335], ASDEX Upgrade [320] and MAST [336]. The newly added data basically reinforces past findings and does not exceed the previously established limit on the product of TPF and halo current fraction. There continues to be significant scatter in both $I_{\text{halo}}/I_{\text{p0}}$ and TPF, which is presumably due to uncontrolled parameters such as the SOL resistivity (which reflects the impurity dynamics during the quench), and/or un-documented differences in MHD activity in the MHD-turbulent post-disruption plasma. The ability to predict (or correlate with experimental parameters), on a first-principle basis, halo current magnitude and TPF remains elusive.

The new MAST and C-Mod data also illustrate some machine-configuration and in-vessel-component sensitivities

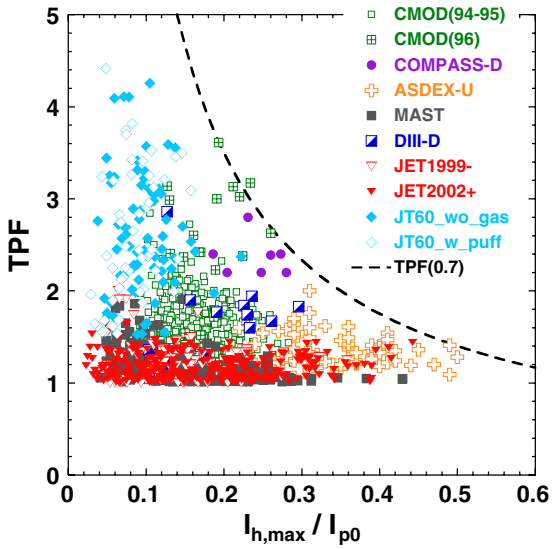

Figure 42. IPB halo current data with recent data from JET and MAST added. The upper bound on the product of $I_{h,\,max}/I_{p0}$ and TPF remains unchanged. Note, however, the lack of new data with TPF > 2 [327].

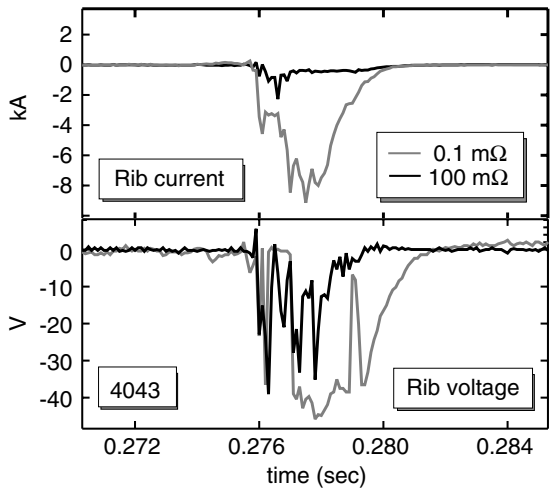

Figure 43. Halo currents and voltages measured in a rib limiter in MAST with low- and high-resistance connection to the vacuum vessel [336].

of the halo current characteristics. On MAST, where several divertor strike-point floor plates ('rib limiters') are connected to vessel ground through instrumented resistors, experiments wherein the resistor value for one of the ribs was varied over a large range (0.1 mΩ, 0.1 Ω, 3.3 kΩ) demonstrated that increasing the resistance resulted in a marked reduction of the halo current flowing through the rib, from a typical value of ∼10 kA down to just a few hundred Amps or less (figure 43). In contrast, the voltage drop varied by no more than a factor of 3 [336]. The halo current flowing through the other, unmodified ribs showed no measurable change. This implies that the disrupting plasma behaves more like a voltage source than a constant current source, and suggests that it may be possible to reduce the halo current in vulnerable in-vessel components by adjusting their resistance to vessel ground. Whether the overall magnitude of the halo current can be reduced by a

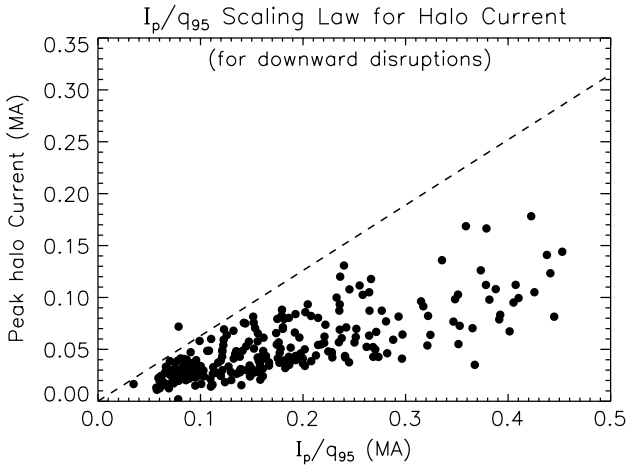

Figure 44. C-Mod halo currents measured in new inboard divertor are reduced by a factor of about 2 compared with the previous scaling (dashed line) with the original inboard divertor.

'global' increase of resistance (isolation) of in-vessel systems from vessel ground remains to be evaluated.

In Alcator C-Mod, installation in 2002 of the new inboard divertor, which has a much less-protruding shape than the previous divertor, has resulted in significant changes in the measured halo current characteristics [337]. The typical halo current scaling versus I_p scaling measured in the new divertor is now about half the magnitude of the scaling measured with the previous divertor (figure 44). During the current quench, the total halo current now seen frequently changes sign in time. The toroidal asymmetries of VDE and non-VDE disruptions (VUDs) are also now quite different, whereas the toroidal asymmetries previously observed were similar. Non-VDE disruptions now have TPF's close to unity. The fact that much of the halo current region (inferred from magnetic reconstruction data) is now observed to miss the new limiter (figure 41) provides a plausible explanation for the reduced halo current magnitude and the transient polarity reversal.

In JT-60U, halo current characteristics have been measured during the runaway current phase [319]. The data show that the in-vessel halo current, averaged toroidally, is very small during the runaway current plateau phase (this absence of halo current is consistent with nearly constant total plasma + runaway current in the plateau phase), but starts to increase after the termination of runaway current, which begins when the plasma surface safety factor, q_s, becomes smaller than 2. Calculation of halo current using a plasma equilibrium analysis code, DINA [338], in which the eddy current in the vacuum vessel is taken into account, shows good agreement with the measured current (figure 45(A)). The toroidal distribution of measured halo current, shown in figure 45(B), indicates that the profile of the halo current has a dominant $n = 1$ toroidal asymmetry with TPF ∼2, and does not change significantly with time. These observations are the same as the previous observations in VDEs without the generation of runaway electrons [332].

Magnetic reconstruction and modelling of axisymmetric halo currents. There has been significant progress since the compilation of the IPB in developing methods that

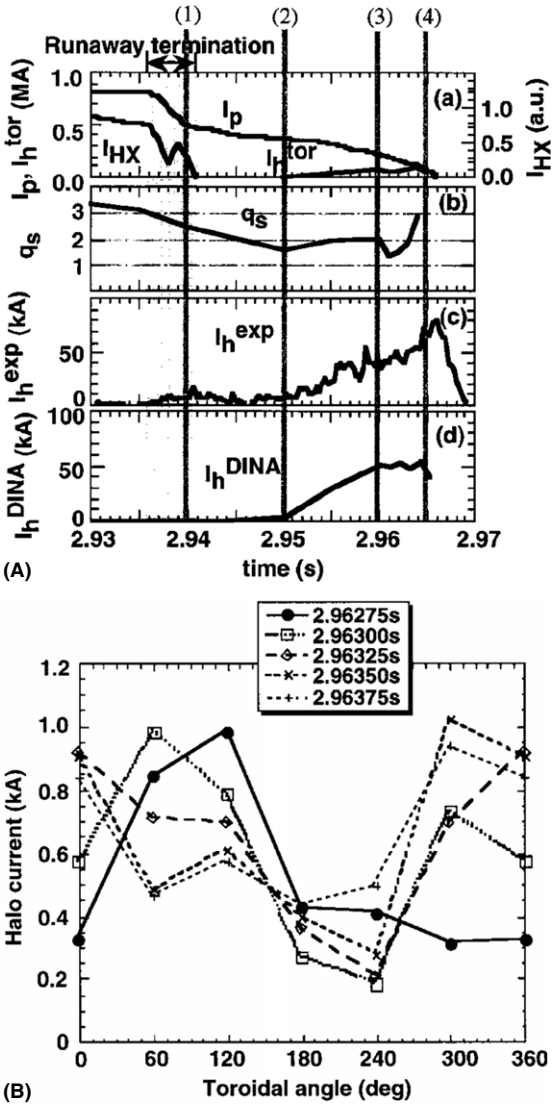

Figure 45. (*A*). Temporal evolution of (a) plasma current, runaway electrons (hard x-ray emission I_{HX}), calculated toroidal halo current I_h^{tor}, (b) surface safety factor q_s, (c) poloidal halo current from Rogowski coil measurements $I_h^{exp}$and (d) that by DINA code analysis I_h^{DINA}, during and after the termination of runaway current [319]. (*B*) Temporal evolution of the toroidal distribution of measured halo current shown every 0.25 ms during the phase of halo current increase (2.962 75–2.963 75 s) [319].

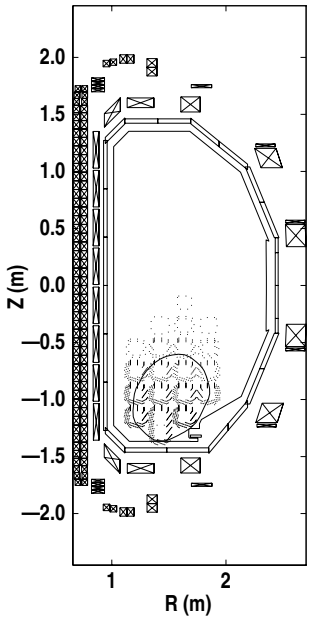

Figure 46. JFIT reconstruction data for the end phase of a DIII-D VDE. The darkness of the contours corresponding to each square current element reflects the amplitude of the current in the element. Elements located in vessel regions showing no contours have negligible current. The solid contour in the plasma region denotes the last closed flux surface, enclosing the 'core' plasma current. Halo current lies outside this core region. Following loss of the last closed flux surface, the halo current region rapidly diffuses to fill the entire vessel. (Reprinted with permission from [331]. Copyright 1999, American Institute of Physics.)

allow magnetic reconstruction of the equilibrium dynamics of the core plasma and halo current regions during a disruption or fast VDE. Traditional equilibrium reconstruction algorithms, based on the accurate Grad–Shafranov (G–S) constraints, frequently have difficulty converging during evolution of VDE's, particularly when large currents are driven in conducting structures. Figure 46 shows a reconstruction—effected using the distributed-current element code JFIT—of the toroidal plasma and halo currents during the final phase of a DIII-D VDE [331]. Since this type of 'current-element-based' code does not impose detailed force balance, but rather fits both conductor and plasma currents, its reconstruction is very robust even in the presence of large driven vessel currents. Robust reconstruction of the post-thermal quench phase of the current evolution allows separation of the core and halo currents and study of their detailed evolution during a disruption. The data in figure 41 are examples of a reconstruction of a C-MOD VUD using a similar current-element method. The alternate method is DINA (G–S equilibrium based) simulation as shown for a JT-60U plasma current termination in figure 45(*A*). Both methods are capable of yielding data on halo current magnitude and width and (by inference) effective halo region conductivity and temperature.

Modelling of core and halo plasma current evolution during DIII-D VDE's and/or the vertical displacement phase of a major disruption has produced improved physics understanding of the driving mechanisms for disruption and halo currents [331]. The halo safety factor is a function of the competing effects of current decay and speed of motion. Rapid core current decay relative to plasma motion (rate of decrease of the core minor radius) tends to increase q_{halo} and thus decrease poloidal halo current (I_{pol}) relative to the toroidal halo current (I_{tor}) (defined as a 'Type I' limit), whereas rapid motion relative to core current decay rate tends to decrease q_{halo} and thus an increase in I_{pol} relative to I_{tor} (defined as a 'Type II' limit). Figure 47 summarizes the measured and model-simulated evolution of a 'Type II' DIII-D VDE . The reconstruction produces a fit of toroidal currents alone, which allows an estimate of the edge safety factor (owing to the virtually complete loss of plasma stored energy producing a force-free plasma). The poloidal halo current can be calculated using the force-free condition in the halo, which implies that $I_{pol} = I_{tor}/q_{halo}$. Following wall contact, toroidal current is induced and convected into the halo from the core plasma. Both toroidal (dI_p/dt) and poloidal ($d\Phi/dt$) voltages

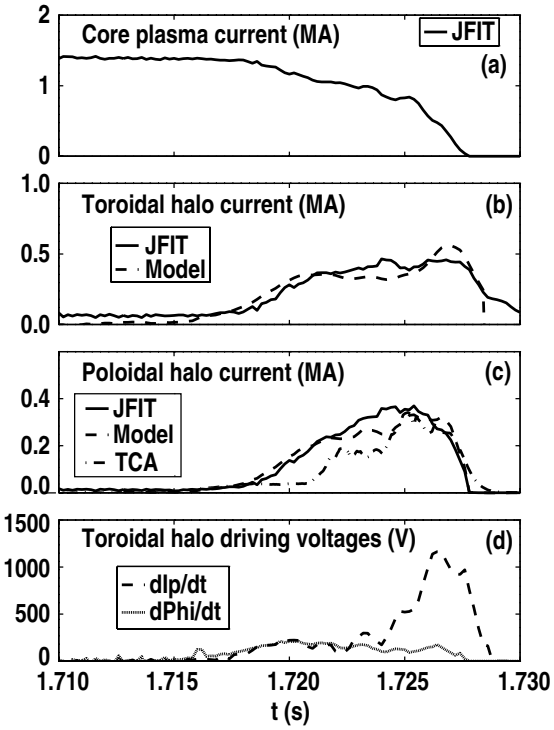

Figure 47. Comparison of measured (TCA = tile current array), modelled and JFIT-inferred halo current for a DIII-D VDE. (Reprinted with permission from [331]. Copyright 1999, American Institute of Physics).

contribute to the halo current drive. The dynamic evolution of this reconstruction-inferred poloidal halo current and the model-simulated current are in good agreement with the value measured by the tile current monitors.

Vertical and radial forces. Vertical and radial forces on the vacuum vessel systems in ASDEX Upgrade and JET are routinely observed during ASDEX Upgrade disruptions and asymmetric VDEs in JET. While the magnitudes and nature of the forces and vessel responses in both cases are commensurate with elementary assessments, the ASDEX Upgrade and JET experiences illustrate some of the underlying complexities and subtleties of the mechanisms responsible.

The ASDEX Upgrade vacuum vessel is equipped with diagnostics to measure mechanical forces and displacements. The vessel is suspended through eight vertical rods (spaced toroidally every 45°), each instrumented with a strain gauge. The vacuum vessel is also instrumented with displacement gauge, which measures its radial and vertical movement at the equatorial midplane at four toroidal locations (every 90°). Over a database of 100 disruptions, vertical forces above 300 kN have been observed with $I_p = 0.8$–1 MA. In slow VDE's (growth rates of order 100 ms), peak magnitudes have reached 500 kN. The peak forces do not scale in a simple manner with the vertical unstable plasma force, $F_{\mathrm{inst,z}} \propto I_p^2$ or , at fixed q_{edge}, with $I_p \times B_t$. The upper limit of the vertical force is described approximately by $|F_z| \leqslant 250$ kN/MA/T. Unlike the large asymmetries seen on JET, the toroidal asymmetries of the vertical forces on ASDEX Upgrade, as measured at the support rods, are typically <20%, and on average ∼10%, and do not show a clear $n = 1$ tilt of the whole vessel. A typical

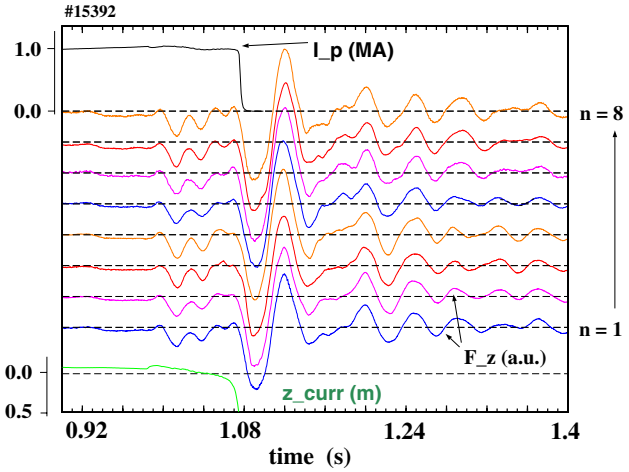

Figure 48. Vertical forces on the ASDEX Upgrade vacuum vessel measured at the eight support rods (labelled $n = 1$–8) [320].

example of the time traces of the eight strain ASDEX Upgrade signals is shown in figure 48. The high degree of azimuthal symmetry is evident.

The displacement of the vessel (measured at the midplane) has a time behaviour similar to the support-rod strain signals. The maximum vertical displacement is $\sim \pm 0.7$ mm at the midplane, while the largest net radial displacement observed was only 0.24 mm, which corresponds to a *static* lateral force of 17 kN. However, since the disruption loads are applied on a timescale much shorter than the vessel oscillation time, the instantaneous lateral forces may be much higher.

A somewhat different type of vessel behaviour during asymmetric VDEs is observed in JET. In JET asymmetric VDEs, the current and position of the plasma can become toroidally non-uniform once the boundary safety factor decreases to a critical value; this is possible when the plasma cross-section shrinks faster than the plasma current decays [339]. The mechanical flexibility of the JET vessel supports, leads to sideways displacements of the vessel. Measurements of these displacements provides clear evidence of the mechanical effects of the forces acting during asymmetric VDEs, against which the predictions of various interpretive models can be benchmarked [340].

Analysis of the interactions among the current-carrying systems (plasma, vessel, TF and PF coils) reveals that the repelling force between the plasma and the vessel is of little significance in asymmetric events, even if it is the main symmetric repelling force. The sideways force acting on the vessel is shown to be mainly due to the interaction between the toroidal field and the asymmetric circulation of currents in the wall [341]. The source/sink of the asymmetric wall current is given by the imbalance between the influx and the outflow of perpendicular (poloidal) halo current at any toroidal location. For a $m = 1/n = 1$ macroscopic asymmetry, the current variation as a function of toroidal angle is expressed by $\delta I_p \cos\phi$. The continuity equation gives the current per unit length flowing vertically from the plasma into the top part of the vessel as $j_v(\phi) = \delta I_p / R \sin\phi$. Apart from a thin ring where the current enters/exits, the wall has zero divergence. Therefore the vessel current density can be expressed in terms

of a stream function and the poloidal components of the wall current estimated as

$$j_s = \frac{\delta I_p}{2R}\frac{p-2s}{p}\sin\phi,$$

where p is the poloidal perimeter of the vessel and s a poloidal co-ordinate (varying between 0 and p). Integration over the vessel of the vertical projection of the poloidal vessel current path gives a net sideways force $\sim \pi B_{tor} b \delta I_p$, where b is the vertical major radius of the torus (full details of this calculation are in [341]).

Studies of plasma current quench dynamics and halo current generation mechanisms and characteristics (magnitude relative to I_{p0} and toroidal peaking factor) indicate that the basic recommendations and design guidelines given in the ITER Physics Basis continue to apply. Some relaxation (increase) in the guideline for the lower bound to the area-normalized current decay time, to $t_{100}/S* \cong 1.7\,\mathrm{ms\,m^{-2}}$, can be recommended, albeit with the caveat that a few examples that fall slightly below this lower bound have now been seen in DIII-D [329]. New data obtained in JET and MAST supports the IPB recommendations for maximum halo current fraction ($I_{h,max}/I_{p0}$) and ($I_{h,max}/I_{p0}$) × TPF product. The new data, including that from Alcator C-Mod, also shows little evidence for large halo current asymmetries, TPF >2. Progress in the development of magnetic reconstruction and simulation methods for assessing plasma evolution and halo region dynamics is providing new data for calibrating predictive models for ITER. Continuing studies of the vertical and radial forces on the ASDEX Upgrade and JET vacuum vessel systems observed, respectively, during disruptions in ASDEX Upgrade and VDEs in JET, demonstrate basic understanding of the magnitude of the vertical loads seen in both cases, but less complete understanding of the radial and/or toroidal asymmetric vertical forces seen in the JET vessel during VDEs with significant toroidal asymmetries.

3.4. Runaway electrons generated by disruptions

The production during disruptions of relativistic (runaway) electrons poses a potentially serious threat to the integrity of ITER plasma-facing-component surfaces. ITER, like any high-current reactor-regime tokamak, will be inherently susceptible to efficient conversion, by Coulomb avalanche multiplication, of plasma current to relativistic (runaway) electron current [1, 289, 290]. Such conversion, of up to 70% of the initial plasma current, is predicted to occur following a naturally occurring disruption, an artificially induced fast plasma shutdown or a loss-of-control vertical displacement event (VDE). The subsequent uncontrolled interaction of this magnitude of runaway current with PFC surfaces has the potential to produce local damage to PFC surfaces and their underlying substrate structures.

Given this inherent susceptibility of ITER to runaway conversion, plus the already-documented instances of significant runaway generation and avalanche multiplication and sometimes even PFC surface damage in the present (TFTR/JET/JT-60U) generation of large tokamaks, there is urgent need to clarify the properties of 'disruption-generated' runaway electrons and to establish methods for avalanche conversion avoidance and runaway discharge termination that can be reliably effected in an ITER-class tokamak. A comprehensive discussion of the physics basis for runaway generation, confinement, loss and potential mitigation has been presented in the 1999 ITER Physics Basis [1]. Since then, significant progress has been made in this subject, and new results and important developments in understanding are summarized in this section.

3.4.1. Observations in present tokamaks. Data continues to be accumulated about parametric sensitivities for runaway generation, confinement and loss in the present generation of tokamaks. Observations reported since the writing of the IPB have been made for example in the JET, JT-60U and Tore Supra tokamaks. Observations documenting runaway generation from injection of argon pellets have also been reported for DIII-D [342]. Various theoretical analyses that contribute to understanding of these results and the corresponding estimated effects in ITER have also been published. Taken together, the new data provide an improved understanding of the nature and possible solutions to the problem of avoiding or mitigating runaway damage in an ITER-class tokamak.

Conditions for runaway generation. Conditions for disruptive runaway generation and/or amplification (avalanche conversion) in present tokamaks continue to receive systematic study. The previously noted tendencies reported in the IPB for disruptive runaway generation to be more prevalent in larger-scale tokamaks and to also be positively correlated with higher toroidal field and/or higher plasma edge safety factor have been clarified by new data. Recent large-tokamak observations of the B_T and edge safety factor domain for disruption runaway electron generation [343, 344] have clarified the findings noted in the IPB. Runaway electrons are observed for $B_T > 2.2\,\mathrm{T}$ and $q_{eff} > 2.5$ in JT-60U, and for $B_T > 2.2\,\mathrm{T}$ and $q_{95} > 2.5$ in JET (typically, in JT-60U, $q_{eff} \approx 1.25 q_{95}$). A similar $\geqslant\sim 2\,\mathrm{T}$ threshold dependence on B_T is also observed in Tore Supra [345]. While definitive understanding of how the B_T threshold for runaway generation scales with plasma size (R and/or a) remains as a future R&D task, these observations and anecdotal data indicating increased tendency towards the disruptive generation of runaways in 'large' versus 'medium' and 'small' tokamaks all suggest that disruption-generated runaways must be expected in ITER.

Direct observation of runaways. Direct observations of disruption-generated runaway electrons in flight have been obtained in JET by application of soft x-ray diagnostics [346]. The resulting SXR data, which dominantly show line radiation of metallic impurities excited by impact of runaway electrons, allow the dynamic evolution of the plasma runaway content to be observed directly. The JET data show that the runaway electrons are generated near the centre of vacuum vessel within a small minor radius, and then subsequently move towards first wall (figure 49). The profile of the runaway beam is observed to be Gaussian and its diameter (FWHM) is determined to be ∼37 cm (compare with an initial thermal plasma diameter of ∼180 cm). Since the local SXR emission is proportional to runaway current density, the q profile can be evaluated: the data indicate $q \sim 0.5$ at the centre of the beam and $q \sim 3$ at the edge. Such information of q profile is essential to examine the beam stability.

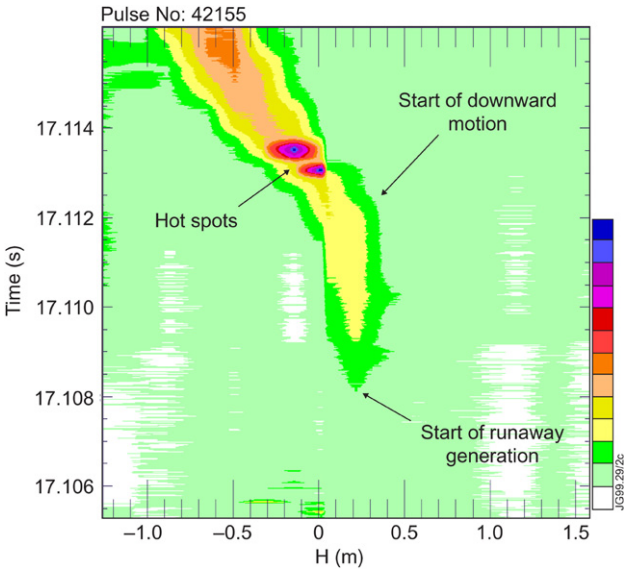

Figure 49. Observation of soft x-ray image of runaways in JET. The downward motion (towards the divertor) is clearly seen. The runaways are first generated 4 ms after the start of the disruption [346].

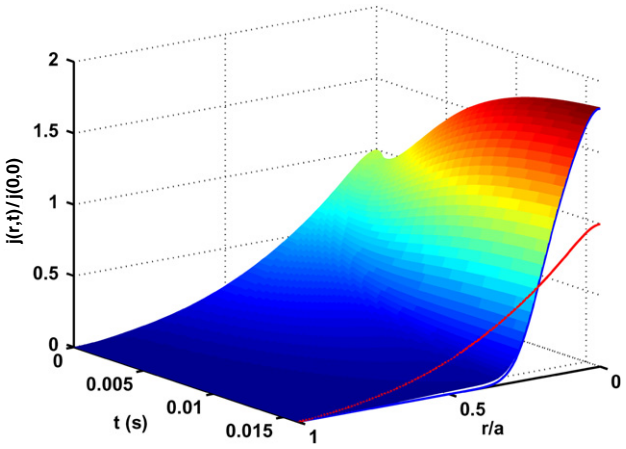

Figure 50. Numerical simulation of the current profile during the runaway conversion phase of a disruption in JET. As is typically observed in experiments, the simulation indicates that the profile of the runaway current is much more peaked than the original current profile (shown by the line at $t = 1$ s) [348].

Theory of runaway dynamics. Progress has been made in calculating the evolution of the runaway population and the electric field during a disruption in a self-consistent way [347–349]. These calculations reproduce the experimentally observed conversion of thermal current to runaway current and thus provide a quantitative link between theory and experiment. In JET, where typically half of the thermal pre-disruption current can be converted into runaways, calculations indicate that much of the final runaway population is already being generated by the avalanche mechanism [344]. Numerical simulations for ITER ($I_p = 15$ MA) employing the same basis suggest that about 70% of the ITER thermal current will be converted to runaway current. These results are in reasonable accord with previous predictions, of up to 80% conversion, cited in the IPB for the 21 MA ITER design.

The present theoretical calculations also show that post-disruption current profile becomes more peaked than the pre-disruption current profile (figure 50) [33, 348]. Previous and present indirect observations of the runaway current channel size and/or profile (e.g. l_i) and the JET SXR imaging data all support this prediction. Peaking has important implications, especially in ITER, for the axisymmetric control and equilibrium stability of the runaway beam. The reason for peaking is that the toroidal electric field diffuses into the centre of the discharge where runaway production is most rapid. This diffusion mechanism is found to be much more important than the one due to toroidal effects [350] mentioned in the IPB. The theoretically predicted peaking probably explains why the central safety factor in JET is $q \sim 0.5$.

Finally, the theoretical calculations suggest that, owing to the high sensitivity of the runaway production mechanisms to local plasma parameters, the runaway current profile can easily become radially filamented. Note the 'hot spots' observed in the JET SXR data in figure 49; these and the other JET observations noted below are likely indications of filamentation.

3.4.2. Interaction of runaway electrons with plasma facing components. The SXR images in JET show that the poloidal width of the wall-interaction (runaway impact) region is relatively narrow: about 10 cm [346]. Observation of discrete pulses (bursts) of hard x-ray emission during the impact period suggests that the runaway current channel itself is filamented. Similar indications (derived from IR TV data) of poloidally localized runaway impact and localized first-wall heating during VDEs are reported for JT-60U [319]. In the case of the JT-60U data, the first-wall heating and HXR indications of runaway loss are attributed to the effects magnetic-fluctuation-produced runaway transport (see section 3.4.3 below). In Tore Supra, local impingement of runaway electrons on the outboard limiter was visible in video images [351].

The runaway filamentation and impact region size observations made in JET and JT-60U support the concerns previously identified in the IPB about the possible PFC-damaging effects of uncontrolled termination of a runaway discharge in ITER. While some uncertainty exists with regard to the exact details of the material damage caused by runaway electron losses, repetitive current quenches with significant conversion to runaway electrons seems unacceptable in ITER. Simulations using the DINA code [352] indicate a runaway electron energy content of ~50 MJ if runaway electrons are confined until $q = 1$ is reached at the edge. However, if runway electron expulsion at $q \sim 2$, as observed in JT-60U [319, 343], is considered, the energy content will be lower (a likely maximum of ~25 MJ). The duration of the runaway electron energy pulse in ITER is estimated to be 130–230 ms depending on the plasma electron temperature during the plasma limiter configuration phase of an uncontrolled VDE. Elementary estimates for ITER based upon 50 MJ runaway energy content, and the measured peak energy deposition of 15–65 MJ m^{-2} [1, 346] give a lower bound on the effective deposition area of ~0.8 m^2. This relatively small area of energy deposition is consistent with expectations due to the small poloidal extent of the runaway electron beam and the imperfect toroidal alignment of affected wall section.

Owing to vertical instability of the runaway current channel, the runaway electron energy will be likely be

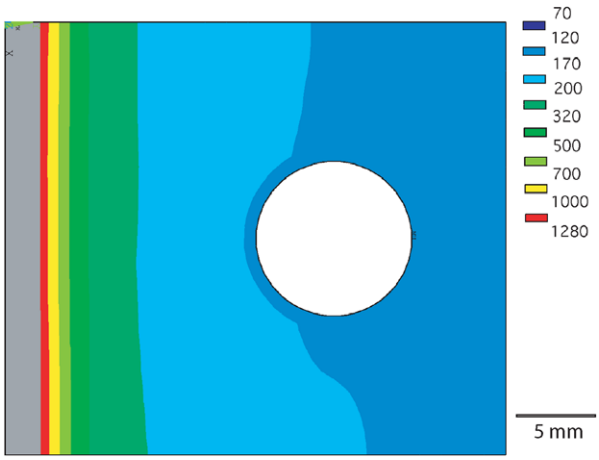

Figure 51. Numerical simulation of the ITER first-wall temperature (in °C), just after energy deposition by 10 MeV runaway electrons, with 50 MJ m^{-2}, deposition time = 0.1 s. From the lefthand side, the simulation geometry comprises 10 mm of beryllium armour, 22 mm of copper heat sink and a 10 mm inner-diameter copper cooling water tube (lateral spacing = 28 mm). The ∼2 mm thick grey zone indicates material attaining temperature larger than the beryllium melting point.(Reprinted from [353]. Copyright 2003, with permission from Elsevier).

deposited on the upper or lower first-wall surface or possibly within the divertor. Hence melting of the beryllium or tungsten armour, or ablation of graphite must be considered. The runaway energy will be deposited on a thin surface layer determined by the electron stopping power and the angle of incidence: the deposition depth is estimated to be ∼2.5 mm for beryllium and carbon and ∼0.2 mm for tungsten. This energy deposition leads to melting of the material in both beryllium and tungsten in the optimistic 15 MJ m^{-2} estimate, and melting plus ablation in the pessimistic 65 MJ m^{-2} estimate.

Similar conclusions are obtained from Monte Carlo code simulations of PFC volumetric energy deposition and temperatures (figure 51) [353]. The simulations suggest that several kg of molten material can be produced (and likely mobilized by gravity and $\boldsymbol{J} \times \boldsymbol{B}$ forces) by a single runaway interaction event. Graphite will undergo some ablation for energy depositions >35 MJ m^{-2}. Again, while significant uncertainties about deposition areas and time scales and effects of runaway beam filamentation and local PFC misalignments exist, frequent occurrence of such an uncontrolled runaway impact event would appear to be unacceptable.

3.4.3. Confinement, termination and mitigation of runaway electrons. The fact that runaway electrons are well confined in an otherwise MHD-stable tokamak discharge, even following disruption, is well known, and many tokamaks observe persistent 'runaway tails' in the after-disruption current decay waveform. In cases where the plasma/runaway current equilibrium is otherwise well controlled, in large tokamaks multi-second post-disruption 'runaway plasmas' can be obtained in this manner. In many cases, the eventual loss of runaways to the PFC surfaces appears to be governed mainly by the gradual (or abrupt) decay of the plasma equilibrium. In direct terms, the runaways are confined until equilibrium decay or control loss 'dumps' them on the PFC surfaces.

Studies of this aspect of runaway confinement and loss, typically at relatively low current (⩽1 MA), have been pursued in JT-60U. The ability of the JT-60U control system to maintain a vertically stable plasma and divertor configuration during the post-disruption phase [354] allows investigation of the subsequent runaway current decay and loss phase. It is found that the post-disruption runaway current tends to decay slowly and smoothly, even with zero loop voltage [343]. It is also confirmed that negative loop voltage increases the decay rate. During the decay phase, it is theoretically found that the observed gradual loss of runaway electrons can be explained by the combined effects of pitch-angle scattering and synchrotron radiation, and this damping can be more efficient than the classical Coulomb-scattering collisional drag [355]. The characteristic decay time of runaway current in JET, 1–2 s, is also explained by these damping mechanisms. In JT-60U, the temporal behaviour of the runaway current can be explained by the balance of avalanche generation and slowing down due to these damping mechanisms [356].

The JT-60U experience demonstrates that when the plasma position and shape control can be maintained after the disruption and formation of a runaway current channel, it is then possible to effect a slow and controlled current termination that can act to benignly terminate runaway electrons and any remaining plasma thermal current.

In JT-60U, it is demonstrated that the runaway electrons that are otherwise well-confined during the after-disruption phase are exhausted at an increased rate owing to the presence of the large-amplitude magnetic fluctuations that spontaneously appear when the edge safety factor q_s becomes small (2 to 3) (figure 52) [319, 357, 358]. The runaway-loss-enhancing effect of magnetic fluctuations, already reported in the IPB, is confirmed by more recent simulations which demonstrate that magnetic islands having widths expected during disruption cause a collisionless loss of relativistic electrons owing to enhancement of the stochasticity of their relativistic motion. This enhancement can explain the enhanced runaway loss experimental observations reported in JT-60U [359] and also the tendencies noted above for runaway 'generation' in large tokamaks to be inhibited at low B_T and/or low-edge q. Low-edge q occurs naturally in a typical VDE and/or the end phase of a VUD; enhanced fluctuation losses seem a likely explanation for the B_T and/or low-q runaway generation thresholds observed in large tokamaks.

In JT-60U, an injection of impurity neon pellets into a post-disruption runaway plasma, caused a prompt exhaust of runaway electrons from the plasma and a reduction in runaway plasma current, without any large-amplitude MHD activity [356]. Since the bulk electron temperature is estimated as very low (∼10 eV), the pellet ablation time would be very long (∼10 s). However, an increment in electron density was observed at the pellet injection. Therefore, it is likely that the runaway electrons themselves play an important role in the pellet ablation, but theoretical and modelling studies to understand this ablation process, together with the exhaust dynamics of runaway electrons, are needed.

Observations shown here were found mainly in JT-60U, where well-confined post-disruption runaway plasmas are stably formed. Such plasmas are suitable to simulate runaway plasmas in ITER and to investigate mitigation techniques

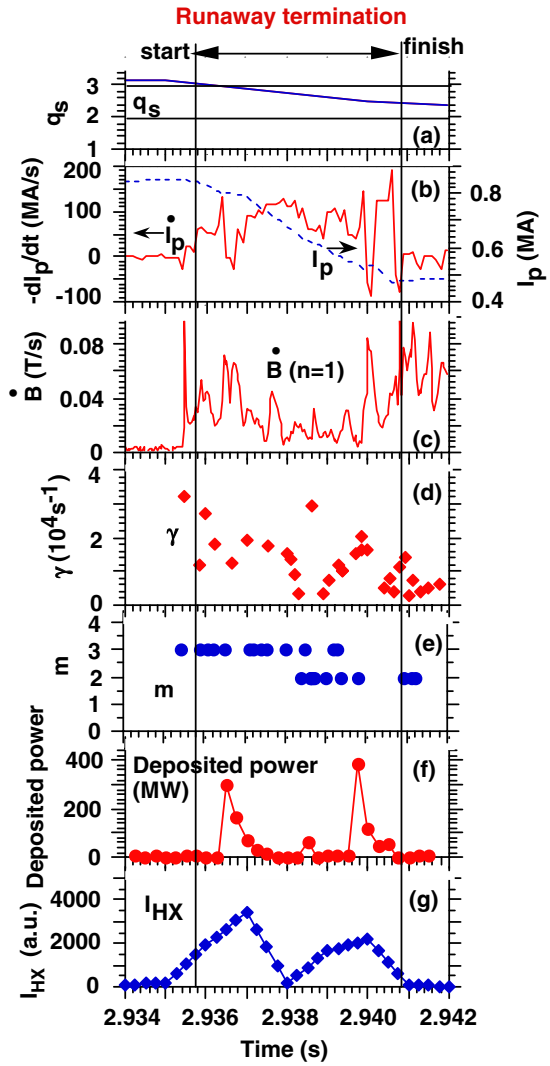

Figure 52. Temporal evolution of (*a*) surface safety factor, (*b*) plasma current and its decay rate, (*c*) magnetic fluctuations with toroidal mode number $n = 1$, (*d*) growth rate γ of each spike in the $n = 1$ mode, (*e*) poloidal mode number m, (*f*) deposited power on the inner divertor plates and (*g*) hard x-ray emission, during runaway termination in JT-60U [319].

for runaway electrons. Recently, stable formation of the post-disruption runaway plasma is reported from FTU [360] and Tore Supra [351]. Their waveforms are very similar to those in JT-60U, which indicates that the physics of the post-disruption runaway plasmas are general among different devices. Further inter-machine experiments and comparisons are strongly encouraged to obtain an even better understanding of the post-disruption runaway plasma.

Fluctuation losses of sufficient magnitude can act to reduce or eliminate avalanche gain. If electrons undergoing runaway acceleration are lost before they have had time to produce knock-on secondaries, they do not contribute to the avalanche growth [361, 362]. This has been shown both numerically [342, 363] and analytically [363]. The exact magnetic fluctuation amplitude necessary to suppress runaway avalanches depends on the mode structure and spectrum of the fluctuations. However, the theory suggests that these dependences are fairly weak and that avalanche suppression should occur, for typical tokamak parameters, when $\delta B_r/B \sim 10^{-3}$, both in ITER and in existing experiments. Further study in present large-tokamak experiments under conditions where appreciable avalanche gain is otherwise evident is required to assess the ability of magnetic fluctuation losses to partially or completely offset avalanche multiplication in a high-avalanche-gain plasma.

The application of an 'equilibrium sustenance and Ohmic ramp-down' method, supplemented by high-Z pellet injection, to terminate ITER runaways has been studied computationally [364]. The studies demonstrate the possibility (subject to PF system voltage, current and power limitations) to effect a controlled ramp down of the currents of the poloidal field coils, in a manner that applies a negative loop voltage (to reduce runaway energy) while simultaneously reducing vertical elongation and maintaining vertical and radial position control. The injection of a succession of small high-Z pellets to enhance radiative and collisional losses during the combined thermal and runaway current ramp down is also proposed. During such a slow termination, there is time to apply additional techniques (e.g. pellet injection [356] or enhancement of natural magnetic fluctuation levels [362]) for the mitigation and termination of runaway electrons. The possibility of mechanical intervention, i.e. a gradual insertion of a movable limiter, also exists.

High-density mitigation. The runaway mitigation scenarios and methods cited above focus on achieving a benign (or acceptable) gradual termination of an already-established runaway discharge. The estimated time scale for effecting such a termination in ITER is likely to be $\geqslant\sim$10 s. An alternate approach of providing a large increase of the plasma electron content, by a factor of $\geqslant$100, before or during the current quench phase of a disruption or VDE, has also been proposed [1]. The basis of this approach lies in the observation that at a sufficiently high electron density (free + bound), the ratio E/E_c becomes less than unity and there is no avalanche multiplication [290]. Here E is the in-plasma toroidal electric field and $E_c = (4\pi e^3 n_e \ln\Lambda/mc^2)$ is the critical electrical field that governs whether or not the avalanche grows. Elementary considerations developed in the IPB show that the corresponding critical electron density, denoted here as the Rosenbluth 'no-avalanche' density, n_{RB}, is given to a reasonable approximation by

$$n_{RB}(10^{20}\,\mathrm{m}^{-3}) \cong 11E(\mathrm{V\,m}^{-1}).$$

Combining this estimate with the further estimate, based upon dissipation of the plasma internal flux, $\mu_0 R l_i/2$, in the minimum current quench time $t_{CQ}/S^* \cong 1.7\,\mathrm{ms\,m}^{-2}$ (see section 3.3.1), leads to the 'no-avalanche' estimates for ITER (given in table 5 of the introduction to section 3) of $E = E_c \cong 38\,\mathrm{V\,m}^{-1}$ and $n_{RB} \cong 4.2\times10^{22}\,\mathrm{m}^{-3}$. As table 5 demonstrates, similar values apply for the present (e.g. JET) and even larger reactor scale tokamaks (e.g. ITER-EDA).

As the IPB discussion notes 'if runaways are to be unconditionally avoided... the electron density must be quite high...' and possibilities of achieving the required ~100-fold increase in density by means of the injection of single or multiple solid pellets and by injection of liquid jets were identified. The likely need, confirmed by impurity injection

simulations [290], for low-Z species injection (H_2/D_2 or He or possible Be) versus medium or high-Z injection to ensure runaway conversion avoidance in ITER-EDA was also noted.

Massive gas injection (MGI) 'disruption mitigation' experiments (see section 3.6.3), effected with subsonic gas jet systems, have demonstrated the feasibility of achieving after-injection electron densities that may have approached the corresponding Rosenbluth no-avalanche density. Experiments with massive ($\sim$350 000 Torr L s^{-1}) He injection in DIII-D [365] and subsequent experiments conducted during 2000–2001 with similarly massive Ne and/or Ar injection ($\sim$100 000 Torr L s^{-1}) [324,366] resulted in attainment of line-averaged n_e approaching 10^{21} m^{-3} and also reduction of the divertor thermal quench loading, peak halo current magnitude and toroidal asymmetry (see section 3.6.3). In JT-60U, similar experiments with the mixed injection comprising a small amount of noble gas (argon, krypton, and/or xenon) and a larger amount of hydrogen gas verified the feasibility of injection schemes with the flexibility to adjust the species mix and ratio of injected gases [367, 368].

In DIII-D, the He and Ne/Ar MGI plasma shutdown experiments resulted in fast plasma current shutdowns that did not produce a detectable indication of runaways. Comparison experiments with Ar pellet injection did produce runaways [342]. Analysis of the He injection data showed that the resulting MGI-initiated current quench plasma exhibited the classical Spitzer resistivity [369, 370]. This finding indicates that ionization and energy balance calculations of T_e and Z_{eff} can be reliably used to calculate the parallel electric field for estimating runaway electron generation and amplification.

Application of the same ionization and energy balance methodology to the 2000–2001 DIII-D experiments by Whyte *et al* suggested that, given the premise that the measured quantity of injected Ar (or Ne) was more-or-less uniformly mixed throughout the $\sim$20 m^3 plasma volume by the end of the current quench, the resulting Ar- (or Ne-) dominated plasma density ($n_{imp} \sim 2 \times 10^{21}$ m^{-3}, equivalent to $n_e \sim$ 2–4 $\times 10^{22}$ m$^{-3} \cong n_{RB}$) would be only weakly ionized, E/E_c would approach unity and the resulting runaway amplification would be small [324, 366]. This could explain the observed lack of runaway production. However, subsequent, further Ar injection experiments, conducted in 2004 with an improved 'high-intensity' gas injection system [371], clearly showed that the injected Ar does not penetrate (in neutral form) more than a few centimetres into the plasma. Subsequent evaluation of the injection flow data also showed that, owing to the finite rise time of the injection flow, the quantity of gas delivered to the plasma by the end of the current quench was significantly less than the total quantity that was injected. Thus the ability of MGI to collisionally suppress runaway electron avalanche remains an area of ongoing study.

The observations reported in [371] also show the central role of MHD fluctuations in effecting mixing and subsequent transport of the edge-ionized neutral gas into the plasma core. Optimization of the injection system to ensure that the desired quantity of gas is delivered to the plasma in a time that is $\leqslant t_{CQ}$ (as short as 4 ms in DIII-D) is also important. These injection system optimization matters are addressed further in section 3.6.3.

In a high-avalanche-gain experiment such as ITER, it is essential that a sufficient injected electron content be present throughout a plasma volume early in the current quench, since this is the critical time to suppress an amplification of seed runaway electrons from whatever source. Adding gas after the formation of the confined runaway electron beam is much less effective with respect to suppressing runaway electrons [372].

3.4.4. Summary and implications for ITER. Experiences with post-disruption runaway generation and parametric explorations of the conditions for runaway generation in the present JET/JT-60U generation of divertor tokamaks, and past experience in the TFTR limiter tokamak, all indicate the likelihood that disruptions, VDEs and even moderately fast plasma shutdowns in ITER will generate high magnitudes of runaway electron current. Observations in JET and theoretical analysis confirms the important role of Coulomb avalanche multiplication, even under present moderate-gain conditions. Extrapolation to the high-gain ITER regime implies a $\sim$70% conversion of plasma current to $\sim$10 MeV runaway current. Natural or uncontrolled rapid loss of this magnitude of runaway current to localized portions of the ITER first-wall or divertor surfaces has a self-evident potential for causing damage.

Experience with runaway confinement in present tokamaks indicates that natural loss processes (those occurring at normal plasma densities and/or with typical pre-disruption levels of MHD fluctuation) will be slow (10–100 s) in ITER. Hence the ability of the ITER plasma equilibrium control system to maintain adequate control during the thermal current to runaway current conversion process and during the subsequent (gradual) runaway termination phase will be critical. Beyond the elementary PF and plasma control system requirements of being able to maintain equilibrium control (see chapter 8 of this issue [373]) of the runaway column, and to then effect a gradual ($\sim$10–100 s) runaway current termination, there also appears the need to take additional action(s)—e.g. by application of a negative loop voltage and/or injection of impurity pellets to exhaust runaways and/or by injection of high-Z gas to promote enhanced collisional losses and/or by introduction of a movable limiter and/or by artificial generation of enhanced MHD fluctuations—to effect a benign runaway shutdown.

Alternatively, injection of massive quantities of neutral particles (and electrons), either prior to onset of disruption (thermal quench) or immediately after thermal quench (within 10 ms after thermal quench onset), theoretically offers a means to unequivocally avoid runaway generation. However, for the technique to be fully successful, plasma electron densities in excess of 10^{22} m^{-3} must be achieved before an appreciable current decay occurs. While present experiences with massive gas injection are modestly encouraging with regard to attaining high electron densities, and while the technology to implement MGI hardware for ITER is straight forward, further study of the underlying neutral penetration, particle/electron assimilation mechanisms and of the role of MHD fluctuations in effecting mixing in an ITER-scale tokamak is required. The possibility that the enhanced level of MHD fluctuations that MGI produces may also act to at least partially offset the predicted Coulomb avalanche gain also needs careful consideration.

Parallel development of more-penetrating neutral particle delivery means, i.e. high-density supersonic gas jets, liquid jets

and/or large-scale pellet injectors with a fast multi-pellet burst capability, would also seem prudent.

Finally, it is important to understand that the susceptibility of ITER (and all presently envisioned after-ITER reactor tokamaks) to runaway conversion and PFC damage owing to uncontrolled runaway current termination arises solely owing to the high-current ($\geqslant$10 MA) nature of these devices. Hence having a 'solution' to ITER's runaway 'problem' is important as soon as current levels approach 10 MA, and having reliable runaway conversion avoidance or mitigation strategies will be essential for a DEMO class tokamak.

3.5. Integrated modelling and simulation

The preceding sections approach the phenomenology, data and physics bases for disruption and disruption effects in a topical manner, as if the phenomenology and effects are separable. The same topical basis for presentation is used in the IPB. However, as the discussion presented therein makes clear, there are cause-and-effect connections among all the observable attributes of disruptions. Thus interpretation of disruption data from present tokamaks and predictions of what will happen in ITER ultimately require the use of some degree of an 'integrated' model. At the very least, the dynamic evolution of the plasma equilibrium configuration and MHD stability—from initial precursor MHD instability growth to final decay of the plasma current—needs to be explicitly taken into consideration. For predictive modelling of ITER disruptions and VDEs, it is also necessary to incorporate or parametrize certain further aspects of the underlying plasma energy balance and current composition (e.g. thermal-to-runaway current conversion in the integrated model). Finally, it is also necessary that models for ITER incorporate an accurate representation of the torus vacuum vessel, in-vessel components and the ex-vessel PF coil system and, in some cases, the before- and after-disruption actions of the PF control system (see chapter 8 of this issue [373]).

Discussion of the physics basis aspects of a complete integrated model can be found in [1]. This subsection discusses recent advances made—in most cases, using 'integrated' dynamic equilibrium models—in the self-consistent numerical simulation of ITER disruptions and VDEs and their resulting EM loading consequences. The use of a variety of integrated models for disruption halo current and runaway data interpretation in present experiments, or simulations of disruption mitigation methods, are described elsewhere in sections 3.4–3.6. Here, representative examples that contribute for reaching the conclusions noted in previous sections about the structural integrity of ITER vacuum vessel and other in-vessel systems are presented. Progress in incorporating more-fundamental 3D MHD instability considerations into these types of simulations is also addressed and a brief summary of the present capabilities and future development needs for integrated modelling is given.

3.5.1. ITER disruption simulations. Prediction of the characteristics and consequences of ITER disruptions must be performed using an integrated simulation model, since the detailed behaviour of the plasma during disruptions and VDEs cannot be directly predicted by simple extrapolation from the experimental 'database(s)' that have been described in preceding sections. The limitation on direct extrapolation is due in part to the different combination of plasma geometry and surrounding passive structures in ITER (as contrasted with geometry and structures in existing machines) and in part to the differences in energy levels and underlying physics processes that arise from the increase in ITER plasma size, current level and energy content.

Much of the integrated modelling or simulation of the effects of ITER disruptions and VDEs has focused on assessments of the basic plasma equilibrium dynamics and the resulting electromagnetic consequences. A self-consistent simulation of the plasma equilibrium dynamics is essential to obtain meaningful estimates and time histories of induced toroidal currents in the vacuum vessel, eddy and halo currents in the in-vessel components and estimates of the location and poloidal extent of plasma–PFC interaction regions. Simulations of the VDE and during-disruption equilibrium dynamics for both present tokamaks and ITER have typically been developed using the DINA [338, 374] and TSC [364, 375] dynamic equilibrium codes. Here, the simulation results for ITER developed on the basis of the DINA code are mainly described.

As the discussion given in section 3.3 details, the plasma current decay time and waveform, and the ensuing dynamics of the vertically unstable plasma configuration evolution, influence the magnitude, location and time history of induced vessel currents and in-vessel-component eddy (circulating) currents. In addition, the current decay rate and equilibrium dynamics influence the magnitude and toroidal asymmetry of the in-vessel halo currents. While it is possible to implement physical models that can calculate core and halo plasma temperature and resistivity (and hence current decay rates) on a first-principle joule-input/radiation-loss basis, a first-principle ability to predict plasma impurity content during disruption is presently lacking. Hence the past and present modelling of ITER disruption and/or VDEs typically proceeds on the empirical basis (for DINA) of adjusting the model parameters to yield the prescribed thermal energy loss and current decay rate or waveform (see figure 53). A similar parametric adjustment procedure (wherein the plasma energy loss and resistivity models are adjusted) can be applied to obtain a prescribed thermal energy loss and/or current decay time or an approximate waveform in TSC modelling. In both the cases, it is important to recognize that, to some extent, the resulting predictions of equilibrium dynamics, halo currents and eddy currents have residual sensitivities to the underlying modelling basis 'input assumptions'.

Since the DINA code (and the TSC code) provide 2D ($n = 0$) plasma equilibrium modelling, only toroidally symmetric halo currents can be evaluated. For the subsequent EM load analyses of the structural effects of in-vessel halo currents, 3D effects (toroidal asymmetries) are introduced in an after-the-fact manner by applying a toroidal peaking factor that is empirically derived from the halo current database (see section 3.3.2) [1, 335].

Two scenarios for disruption are simulated.

- Major disruption (MD), which starts from a thermal quench and q profile flattening when plasma is at its nominal position. This is followed by a fast reduction of

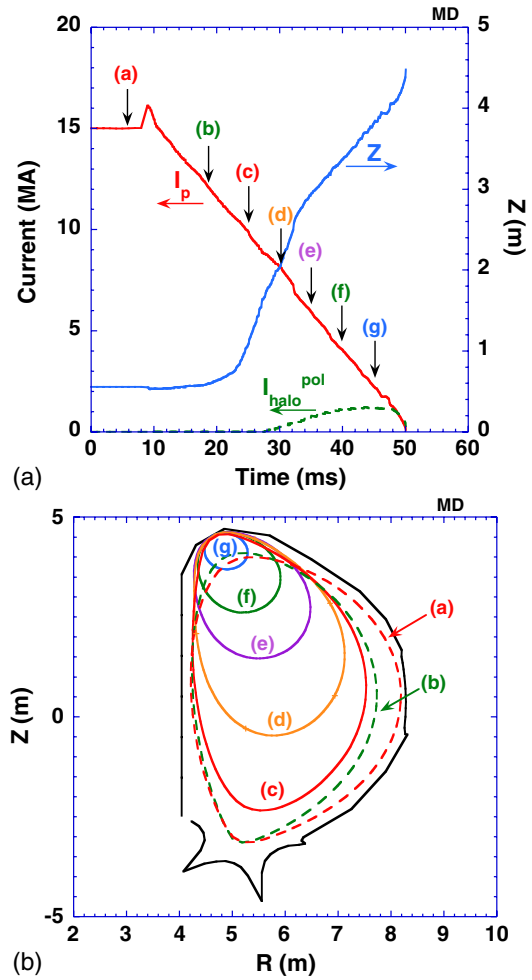

Figure 53. Time evolution of (*a*) plasma and poloidal halo currents, (*b*) plasma boundary (LCFS), for the MD case in table 9.

plasma current. In this case, the ex-plasma halo begins to form immediately following the end of the thermal quench.

- Vertical displacement event (VDE), which starts from a vertical displacement due to loss of vertical position control. Both downward and upward plasma vertical movements are considered. After the plasma boundary touches the wall, the value of q at the plasma boundary decreases as the plasma cross-section area diminishes. When the boundary q reaches a specified value q= 1.5–2.0 [327, 376], thermal quench takes place triggering a fast reduction of the plasma current. For this case of disruption, the halo area shows up just after the formation of a limiter configuration.

The initial plasma equilibrium is taken from the 15 MA ITER 'reference' inductive scenario ($l_i = 0.85$, $\beta_p = 0.7$). Major assumptions used in the two simulation cases are summarized in table 9.

3.5.2. DINA disruption modelling basis. For disruption and VDE modelling, an extended 'disruption simulation' version of the basic DINA code has been developed. The extensions provide the ability to simulate the dynamic evolution of 2D plasma equilibrium with both closed (core) and open (SOL) magnetic surfaces. The plasma equilibrium modelling is achieved in combination with an axisymmetric model of the external circuit (PF coils and surrounding conducting vacuum vessel and in-vessel structures). Flux-surface-averaged plasma energy transport equations are also solved simultaneously. The resulting integrated model has been validated in several tokamaks [377, 378].

Table 9. Major assumptions used in ITER disruption simulations.

Parameter	MD	VDE
Current quench rate (time)	375 MA s^{-1} (40 ms)	375 MA s^{-1} (40 ms)
Current quench waveform	Linear	Linear
Thermal quench duration	1 ms	1 ms
Beta drop during vertical movement (before thermal quench)	0	≈0.3
Surface q value at thermal quench	3	1.5–2
Decrease of l_i during thermal quench	0.15–0.2	0.15–0.2
$T_{e,\,halo}/T_{e,\,core}$	1	1

The DINA disruption simulation incorporates a detailed axisymmetric representation of the ITER vacuum vessel and in-vessel systems. The vacuum vessel is modelled by a set of thin plates with relevant resistance, so that the global L/R time can be matched with that calculated for the actual geometry. Blanket modules are modelled by a set of axisymmetric current-carrying rings, with opposite rings connected in the toroidal direction so that a zero net toroidal current flow is obtained. Resistances of each pair of rings are selected to yield an equivalent penetration time for poloidal magnetic fields through a blanket that is calibrated with a 3D analysis of eddy current dynamics obtained from finite-element simulations.

During thermal quench, conservation of helicity is invoked to calculate the resulting redistribution of the plasma current profile. After-thermal-quench energy transport is not calculated and both plasma core and halo region electron temperatures are assumed to be uniform, with a prescribed value. A temperature ratio for the plasma and halo areas can be specified (see table 9). During the current quench phase, the width of the halo region is determined by a simple model, benchmarked against JT-60U disruption data [379], that is based on conservation of toroidal flux within the plasma cross-section and which is taken to include the halo region once vertical instability develops. Application of this flux-conservation model to ITER yields a halo region evolution that mimics what is seen in JT-60U, wherein the halo spatial width is observed to increase gradually as the plasma moves downwards during a VUD.

A somewhat different type of halo evolution is observed in DIII-D, where observations and simulations indicate that the halo spatial width stays approximately constant during the course of a VUD [331, 369]. While the difference between the two classes of halo-width observations is not large, and while it is likely that both types of observations are consistent with a common flux-conservation and diffusion physics basis model, further physics basis understanding and development of a more-predictive halo model is required. Such a model should incorporate internally self-consistent temperature and conductivity profiles for the halo region, to allow refined halo width and halo current magnitude predictions for ITER.

3.5.3. Predictions for major disruptions and VDEs in ITER. Since the impact of vertical displacements on the machine, due to eddy and halo currents, is more severe for downward displacement (towards the X-point) than for upward displacement, the initial plasma position in ITER is specified to make the plasma movement upwards. The initial position, where the plasma is balanced and will move randomly upwards or downwards at disruption, is known as the neutral point [334]. The neutral point of up–down asymmetric, single-null diverted plasmas, however, depends on the change of l_i and β_p at thermal quench [330], implying that no unique neutral point exists. Similarly the present initial plasma position in ITER is chosen to ensure that the upward movement is obtained for a wide range of l_i and β_p changes. This has been confirmed by systematic simulation studies for a range of MD cases, and for example, figure 53 shows the time evolution of the plasma and poloidal halo currents and the plasma boundary (last closed flux surface) for an MD case simulated by the DINA code. It is found that for the changes of l_i and β_p at thermal quench prescribed in table 9, the plasma moves upwards after the thermal quench. The maximum in-vessel (poloidal) halo current reaches about 1.5 MA ($\sim$10% of I_{p0}).

Figure 54 shows the time evolution of the plasma and poloidal halo currents, plasma vertical position and (b) plasma core (LCFS) and halo boundaries for a downward VDE case. A small initial downward perturbation is provided at $t = 0$, and the plasma moves downwards, without vertical control, before the thermal quench occurs at $t = 670$ ms. The peak halo current for this VDE reaches about 3 MA ($\sim$20% I_{p0}).

3.5.4. Calculation of non-axisymmetric halo current effects in ITER by M3D code. Full three-dimensional simulations of VDEs in ITER have been carried out using the M3D code [380, 381]. Scaling of the VDE vertical growth rate proportional to wall resistivity has been verified. Simulations have been done of disruptions caused by large-inversion-radius internal kink modes, as well as by the non-linear growth of resistive wall modes.

The M3D code includes resistive wall boundary conditions, which match the solution inside the resistive wall to the exterior vacuum solution. The exterior problem is solved with a Green's function method, using the GRIN code [382]. The M3D code has a time-dependent self-consistent resistivity model and realistic ITER geometry including the magnetic separatrix. Thermal conductivity in the presence of the separatrix provides an adequate temperature contrast between the core and the halo region (the open field line region outside the separatrix).

The M3D simulation capabilities allow modelling of both the disruption and VDE phases of a VUD, where disruption causes a thermal quench, which in turn causes a current quench, vertical instability and halo current development. The example in figure 55 shows a VUD initiated by an unstable large-inversion-radius internal kink mode. When the internal kink instability becomes sufficiently non-linear, toroidal coupling to other modes causes disruption. The plasma cools because of transport along stochastic field lines. This raises the resistivity and dissipates the current. The current decay is accompanied by a VDE and halo current development. The VDE growth rate is at least twice as fast in the presence of disruption. The

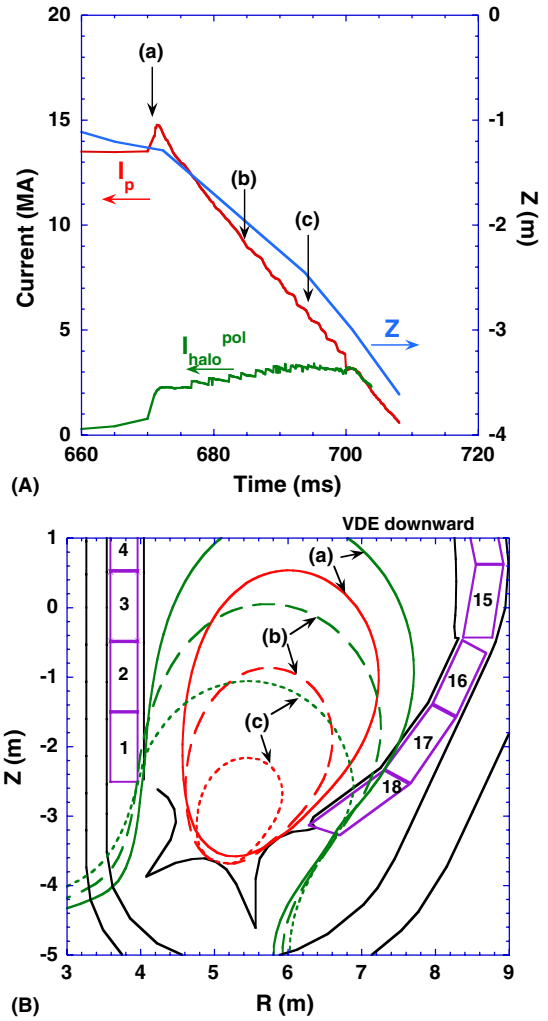

Figure 54. Time evolutions of (A) plasma and poloidal halo currents, vertical position (Z) and (B) plasma LCFS (closed ellipses labelled '(a)–(c)') and halo boundaries (open curves labelled '(a)–(c)') for the downward VDE case [327].

calculated toroidal peaking factor (TPF) for the poloidal halo current can be as large as 3 but typically tends to about 2. This range of TPF matches the observations for present experiments in the halo current database (see section 3.3.2).

Simulations have also been made of disruptions caused by resistive wall modes, which, in turn, are triggered by the application of static non-axisymmetric magnetic error fields. A parallel viscosity model has been introduced in M3D to provide the dissipative coupling needed to reproduce resistive wall mode physics. The model conserves toroidal angular momentum and is consistent with equilibria with toroidal flow in which the equilibrium angular frequency is a flux function. When the external magnetic perturbations are applied, the parallel viscosity causes a damping of the toroidal flow, growth of the plasma RWM and onset of disruption.

The example simulations presented here illustrate the potential of the M3D model, with incorporation of appropriate 'auxiliary' MHD and energy transport physics models, to simulate the full sequence of events for ITER-like VUDs and also hot-plasma VDEs. Validation of the overall simulation capabilities and/or auxiliary models, by comparison with

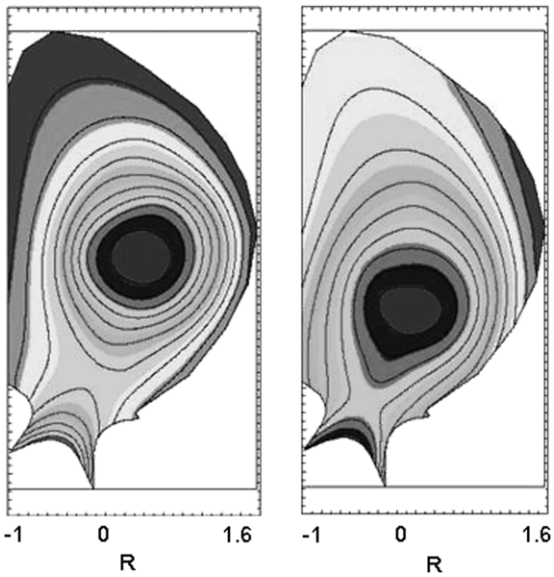

Figure 55. An example of M3D calculation for a VUD caused by a large-inversion-radius internal kink mode. Contours of poloidal flux at two stages of the VDE development are shown [383].

specific VUD and VDE data from present experiments, remains a future R&D task.

3.5.5. Summary and future needs. The development and application of dynamic 2D plasma equilibrium models for internally self-consistent disruption and VDE data interpretation, and simulation of ITER disruption and VDE scenarios, has reached a reasonably mature state. Modelling/simulation accuracies (self-consistency) seem adequate for purposes of providing ITER vacuum vessel and in-vessel component electromagnetic loading data. While the basic equilibrium evolution dynamics aspects of the presently-available models have been well validated, validation of some of the other modelling elements (e.g. halo current dynamics and core and halo region energy balance) is not as complete or universal. Hence predictions from 2D models must be accompanied by due recognition of the ranges of outcomes that uncertainties in the underlying disruption physics basis, or auxiliary modelling element basis, introduce. Needs for sensitivity studies to investigate the range of outcomes and for improvements upon the present empirical or ad hoc nature of some of the auxiliary modelling bases are also obvious.

The emerging development of full 3D MHD dynamic equilibrium modelling capabilities is promising and offers prospects for a 'fully self-consistent simulation' of ITER disruption scenarios and consequences. The application of 3D dynamic modelling to the interpretation of disruption data in present tokamaks will likely aid in sorting out some of the presently obscure internal cause-and-effect dynamics of the precursor-growth to thermal-quench phase of disruptions and perhaps also help in clarifying the cause(s) for the wide range of data scatter in current quench and halo current databases. Finally, as for the 2D modelling, incorporation of adequate (and validated) auxiliary modelling elements into the 3D models will be needed before an 'ITER disruption simulation' can contain sufficiently realistic physics to make quantitative predictions.

3.6. Disruption avoidance, prediction and mitigation

The importance for ITER of avoiding, wherever possible, the occurrence of disruption or loss-of-equilibrium-control events that lead to disruption is universally acknowledged. Similarly, the need for mitigation strategies that minimize or eliminate the adverse consequences of disruptions is well known. Finally, there are operational reasons—mostly related to protecting the integrity and continued function of PFC components—why a fast plasma power shutdown capability is needed. The combination of these considerations leads to a compelling need to develop disruption prediction, avoidance and mitigation strategies that can be effected with high reliability, while at the same time allowing ITER scientists to conduct exploratory plasma development and optimization studies in burning plasma regimes where the risk of disruption will be appreciable.

The needs and considerations noted above have been identified and discussed in the ITER Physics Basis [1]. In the three sections that follow, recent progress in (i) disruption prediction (section 3.6.1), (ii) disruption avoidance (section 3.6.2) and (iii) disruption mitigation and fast plasma shutdown (section 3.6.3) are summarized. There is overlap among these topics and some of the physics basis findings already cited in the previous sections of this chapter. Overall implications for ITER and needs for further research are summarized in the conclusion of section 3.6.3.

3.6.1. Disruption prediction. Having a 'real-time' capability to predict impending disruption is acknowledged as critical to being able to operate a tokamak with adequate disruption protection (to the extent that subsequent resumption of plasma operation in a timely manner will not be compromised). Here 'real-time' denotes the basing of disruption prediction on real-time plasma and device diagnostic data, to provide sufficient advance indication of impending disruption so that avoidance and/or pre-emptive mitigation measures can be taken. The specifics of what constitutes 'sufficient' depends on the avoidance and/or mitigation techniques to be employed and also on the inherent time scale of the tokamak in question. These specifics and requirements for real-time prediction in present tokamaks and ITER are discussed below.

The implementation of a reliable real-time disruption prediction or an advance warning method is a challenging task that has been investigated in various tokamaks over the last decade. In many experiments, 'deterministic' prediction (warning) methods based upon monitoring of one or two key indications of impending disruption—for example, locking of an initially rotating tearing mode and/or bolometric indication of a global plasma energy balance deficit—are deployed and in some instances used in a real-time manner to initiate remedial or pre-emptive actions (see section 3.6.2). While such 'disruption avoidance' methods are often effective for their intended application, they do not always reliably predict the impending occurrence of disruption over a wide range of plasma parameters or for all phases of a discharge. Such methods may also produce undesirable 'false-alarm' indications of pending disruption in cases where the plasma would not actually disrupt. In situations where alarm indications are coupled to pre-emptive action, too-frequent

false alarms can thwart progress in the development of high-performance plasma operation regimes.

To address some of the perceived shortcomings of deterministic prediction, the neural network (NN) technique—which draws upon a wider range of input data to develop a single disruption predictor output—has been applied. A significant NN development and test effort has been implemented on present experiments, including ADITYA, ASDEX Upgrade, DIII-D, JET, JT-60U and TEXT. Performance of a NN is typically quantified in terms of the success rate, SR, the fraction of instances where the NN successfully predicts the occurrence of disruption or, alternately, the failure rate (FR) or the missed alarm (MA) rate (FR $=$ 1 $-$ SR), wherein the NN fails to detect an impending disruption. The false alarm (FA) rate, where the NN indication of pending disruption proves to be incorrect, is a second important performance attribute.

Neutral networks require 'training': before-deployment development of the weighting coefficients that control the summation of the input signals. This training requires a multi-discharge data set—time histories of the relevant input parameters—for discharges that disrupt. The training is specific to the NN input data set being utilized and also, at least to some extent, to the plasma operation mode and operation attributes of the particular tokamak that the NN is being implemented on. Once the NN is trained, its performance can be validated 'off-line' against a second (independent) already-archived data set or tested, either 'off-line' or 'on-line', against newly-acquired or real-time data.

NN development work performed at DIII-D used a NN to predict the value of the β_N at the disruption starting from a 33-input set of data acquired during NBI-heated discharges [384]. The NN was trained to predict, some tens of milliseconds in advance, the maximum β_N that a given plasma would reach at the time of disruption. In an off-line validation test, the NN had a $\geqslant$90% SR in accurately predicting β_N at disruption, but the FA rate was 20% for non-disruptive cases. The same type of 'simple' NN method was tried to predict the density limit disruption in the ADITYA Tokamak [385].

In ASDEX Upgrade, a NN predictor was used to trigger killer pellet injection as a mitigation action for density limit disruptions [386]. The NN output, developed from 13 input parameters sampled every 2.5 ms, predicted the time interval, Δt_{NN}, before disruption. An alarm action was activated for $\Delta t_{NN} \leqslant 50$ ms. A database from 99 disruptive shots and 386 non-disruptive shots was used for training. In an off-line (open-loop, without pellet injection) validation, an 85% SR (55 of 65 disruptive shots in the validation set) was achieved, with a 1% FA rate for 500 non-disruptive shots. Figure 56 shows an example of open-loop alarm indication. The disruption is predicted more than 20 ms before occurrence.

In JT-60U, a NN-determined 'stability level', calculated from 9 input parameters, sampled every 2 ms, is used to predict the occurrence of disruption [387]. The NN was trained in two steps, first with 12 disruptive and 6 non-disruptive shots (Step 1) and second (Step 2), with the NN output data for 12 disruptive shots modified according to the output stability levels from the NN trained in Step 1. The resulting optimized 'neural net disruption predictor' was tested against 300 disruptive shots and 1008 non-disruptive shots selected

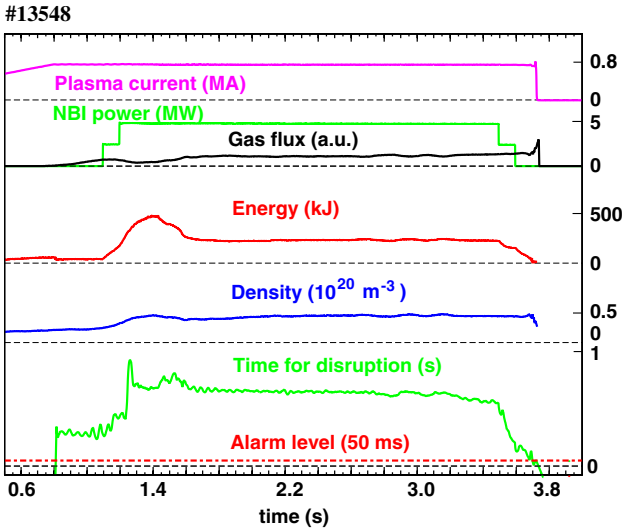

Figure 56. Neural network prediction for an ASDEX Upgrade density limit disruption that follows NBI turn-off [386].

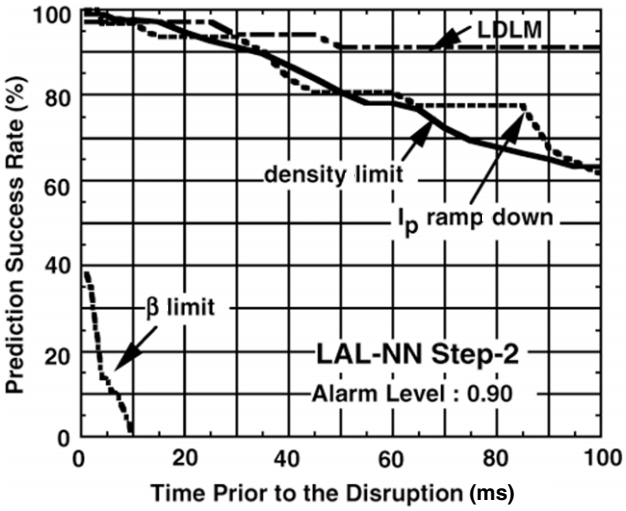

Figure 57. JT-60U: Prediction success rates of the major disruption caused by the density limit (206 shots), the plasma current ramp down (31 shots), the low density locked modes (LDLM) (34 shots) and the β-limit (29 shots) [387].

from 9 years of JT-60U operation. The SR for predicting disruptions—other than those caused by density limit, high l_i during current ramp down or locked modes at low density—with 10 ms advance warning was 97–98%. The FA rate for non-disruptive shots was 2.1%. Careful selection of input parameters and the two-step training method reduced the FA rate, resulting in a considerable improvement of the overall prediction success rate.

The effects of the type (cause) of disruption on the SR are shown in figure 57. For low-density locked mode disruptions, the SR is a satisfactory 100 ms before disruption, despite the fact that the locked mode data was not included in the training data set. In contrast, the prediction of the β_N limit disruption has a much lower SR, probably because there are no clear precursor indications in the input set that occurs more than 20 ms before disruption. Recent investigations in JT-60U have shown that training with non-disruptive shots and careful

Table 10. JET NN performances. Δt_p is the time before the disruption. The numbers in brackets indicate the total number of the shots for training, validation and performance test [389].

	Δt_p (ms)								
	40	60	80	100	120	140	160	180	200
Training set									
MA (86)	2	3	4	3	4	5	8	9	17
FA (400)	0	0	0	0	0	0	0	0	0
Validation set									
MA (35)	0	0	1	1	2	3	4	5	5
FA (246)	0	0	0	1	1	1	1	1	2
Test set									
MA (62)	9	11	13	10	13	17	18	21	22
FA (132)	0	0	0	0	0	0	0	0	0

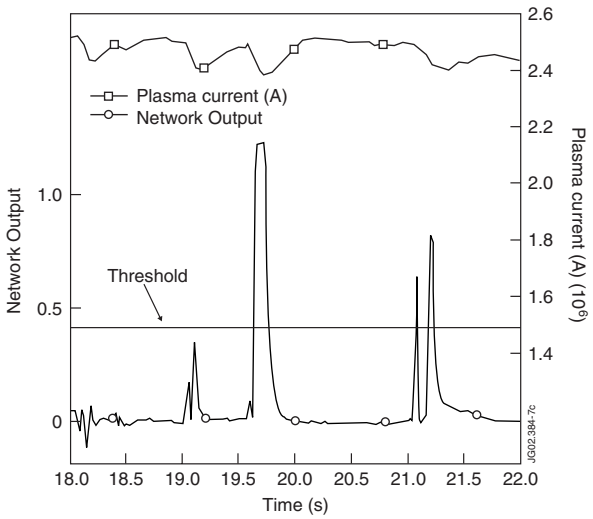

Figure 58. False alarm generated in JET shot 46314. Alarm indication occurs at 19.7, 21.1 and 21.2 s [389].

adjustment of the Step 2 output (β-limit) data can improve the beta limit prediction and reduce the FA rate [388].

Various NN optimizations have been applied for the prediction of disruptions in JET [389]. The data sampling interval is 20 ms. The performances of the best NN optimizations are shown in table 10. The FA rates are negligible for both the validation and the test sets over the 20–200 ms time window considered. The number of MAs is larger but comprises approximately 11% of the entire validation and test data set (97 shots) at 100 ms before the disruption.

For the best-performing NN configuration considered by Cannas *et al* the MA rates are 10–20% (SR = 80–90%) for the test data set, whereas the FA is low (zero). This balance of the SR versus FA rate arises owing to the penalty applied, during the NN training process, to the occurrence of a FA, which, in the intended application, would invoke a 'soft landing' of the plasma current (see section 3.6.2), thereby preventing completion of the experiment. A different choice in the penalty of the NN output signal would achieve a different balance between the SR and FA performance.

For JET disruptions caused by tearing instability or the density limit, the NN prediction provides satisfactory reliability and a low FA rate for look-ahead times up to a few hundreds of milliseconds before the disruption occurrence. A similarly satisfactory result is obtained in ASDEX Upgrade and JT-60U. Better results with regard to high SR are obtained if the operations-regime and disruption-type space parameters used for the training and validation of the NN are well covered and representative of the parameters that are encountered upon NN deployment. On the other hand, further NN development is necessary to reduce the number of FAs generated: in real-time deployment NNs, unexpected FAs (figure 58) are sometimes encountered.

Discussion and future work. The experience gained with NN disruption prediction methods has identified a number of issues that call, on the one hand, for a dedicated effort in the generation of the input signals and, on the other hand, for a deeper understanding of disruption dynamics and precursor identification.

Regarding the first aspect, any kind of disruption prediction means, NN or other, requires a dedicated set of input signals with high reliability, adequate time resolution and long-term calibration stability. Signals used to date for NN prediction have typically been generated for other purposes and have some shortcomings for NN testing and deployment. In addition to 'quality' considerations, it is essential that NN input data include signals that provide precursor indication of pending disruption. For some classes and/or overt causes of disruption (e.g. density limit or cold-edge disruption, see section 3.1), ample precursor indication is present in multiple signals. For other cases (e.g. fast-growing internal MHD triggered β-limit or pressure-gradient-limit disruptions), signals with sufficient precursor content are more difficult to obtain.

There are also other critical points that are intrinsic to the NN approach. These points include limitations on extrapolation capability, the need during initial training for a sufficiently large and comprehensive database of disruptive discharges and the machine-specific nature of the NN design. The extrapolation capability of NNs is intrinsically poor in situations where expanded or new operation regimes are encountered. In present experiments, plasma operation parameters naturally evolve with hardware improvement and new abilities to produce plasmas with higher performances. Hence the training of NNs requires updating as the plasma operating domain and performance capabilities are expanded. This requires having a database that covers the full operational parameter space of the machine, as was tried for example in JT-60U for 13 years of operation [387]. Furthermore, for developing a machine-independent disruption prediction method applicable to a next-step device, a more comprehensive database that covers a variety of past and presently-operating tokamaks in a more 'dimensionless' manner appears essential. Finally, given the data quality and stability concerns noted above, quality and calibration stability of the data that is to comprise this multi-machine database will need to be given careful consideration.

Recently, a first-of-a-kind cross-machine prediction of disruption occurrence has been undertaken, using data from JET and ASDEX Upgrade, with a NN predictor trained

on seven normalized dimensionless plasma parameters plus normalized time. The NN was trained with JET data and tested on ASDEX Upgrade data, with a resulting SR of 67%; conversely a NN trained with ASDEX Upgrade data and tested on JET data showed a 69% SR [390]. These initial inter-machine NN development results are an encouraging first step and indicate a possible way to overcome some of the concerns about NN portability and the need to have data from disruptive operation of a future device (ITER) before the NN predictor for disruption-onset prediction can be developed.

Another realization that has emerged from present studies is that an integrated set of NNs, each optimized to predict a given 'type' (cause) of disruption, will probably be required. Hence classification and characterization of observable precursors and onset processes for each 'type' of disruption will be needed (see section 3.1), and during-operation refinement of ITER NN predictions will undoubtedly be required. Nevertheless, in a next-step (after ITER) reactor tokamak, where the plasma configuration and operating mode will necessarily be well defined and not subject to major evolution, neural net prediction of impending disruption will be a promising, but not exclusive, means available for activating disruption avoidance (and ultimately mitigation) schemes.

3.6.2. Disruption avoidance. The number of disruptions in ITER must be kept as low as possible. In principle, disruptions can be avoided during operation by an accurate construction of the machine itself, careful selection of the plasma operating space and discharge development procedures (here 'discharge development' denotes the plasma 'start-up' procedure that produces the final stationary plasma magnetic and kinetic equilibrium state) and deployment of a sophisticated plasma control system (see chapter 8 in this issue [373]). Careful alignment of the magnetic coils, for example, will reduce the tendency of MHD modes to lock owing to the presence of finite non-axisymmetric magnetic field errors [1, 391, 392]. Precise adjustment of first-wall elements will reduce the probability of local PFC surface overheating and impurity production. Precise control of the magnetic and kinetic evolution of the plasma configuration will prevent unwanted wall-interaction events and help ensure that the plasma development and equilibrium state successfully avoids the various 'causes' of disruption that apply. Finally, given the need for reliable and effective plasma control, it is self-evident that to effect disruption avoidance, all the operation systems involved (hardware and software) must be highly reliable and fail-safe.

The selection of ITER-applicable operation scenarios with low-disruption rates is a further and inherently problematic issue. Burning plasma operation in ITER will require operation in regimes with small margins against each of the three major plasma operation limits (see section 3.1) and under conditions where the need for external stabilization of NTM and/or RWM MHD instabilities and active control of divertor attachment (power flux), plasma impurity content and fusion power is anticipated. So achieving low disruption rates in ITER will require a combination of precise and reliable plasma control, reliable operation of external stabilization systems and careful development of plasma scenarios that successfully skirt the multiple 'causes' of disruption that lie along the plasma 'start-up' trajectory.

This section continues the discussion of disruption avoidance found in the ITER Physics Basis [1]. Examples of experimentally implemented disruption avoidance techniques will be presented. As the discussion presented in section 3.6.1 tries to make clear, the emphasis here is on the 'passive' or 'single-parameter' deterministic and active control aspects of disruption avoidance, rather than on the 'multi-parameter' prediction means that can also provide indication of pending disruption.

Disruption avoidance techniques employed in present tokamaks (and envisioned for ITER) are divided into two categories:

(i) machine preparation and design of 'safe-pulse' schemes and
(ii) real-time intervention with single or multiple feedback control loops or algorithms—often based upon physical models—to control specific plasma regimes and/or recover from failure states [393]. Recovery procedures can include effecting 'soft landing' or 'soft-stop' measures that attempt to end the plasma discharge without producing a full-current or full-thermal-energy disruption.

Faster-acting plasma shutdown and machine protection processes (sometimes called 'hard stop' action(s)) intended to be applied in situations where disruption appears to be otherwise unavoidable are separately discussed in section 3.6.3.

Passive avoidance. The need for thorough and consistent pre-pulse conditioning of the plasma-facing wall (by dc or HF glow discharge cleaning [394]) and by the periodic deposition of first-wall coating materials (e.g. boron) is well known. Wall conditioning has a distinct impact on the plasma impurity content and therefore on the density limit. Wall conditions also affect the plasma breakdown and initial current channel development phase of the discharge and hence 'start-up' reliability (see chapter 8 of this issue [373]).

The importance of avoiding mode locking during the plasma start-up and initial low-density Ohmic flattop phase is also well known. Disruption-avoiding 'pulse designs' that prevent mode locking during pulse ramp up have been demonstrated on TCV [395]. The operation space for locking of MHD modes was statistically investigated in terms of plasma density, safety factor, shaping (elongation and triangularity) and vertical position. A clear tendency to develop mode locking is found for low density and low plasma shaping. Modification of the shaping in the pulse ramp-up phase successfully avoids mode locking (figure 59). Findings [396] about the need for low levels of non-axisymmetric error field for locked mode avoidance in low-density DIII-D Ohmic plasmas and in higher-density high-β NBI-heated plasmas [392] have already been reported in the IPB.

Disruption avoidance strategies. In present experiments, so-called 'soft-landing' or 'soft-stop' procedures comprise the most-frequently applied action to avoid occurrence of impending disruption. Here, following detection of an impending disruption (see previous section) or violation of one or more of various plasma operation 'monitor limits', the plasma current is immediately ramped down in an externally controlled pre-defined way [393]. This action does not

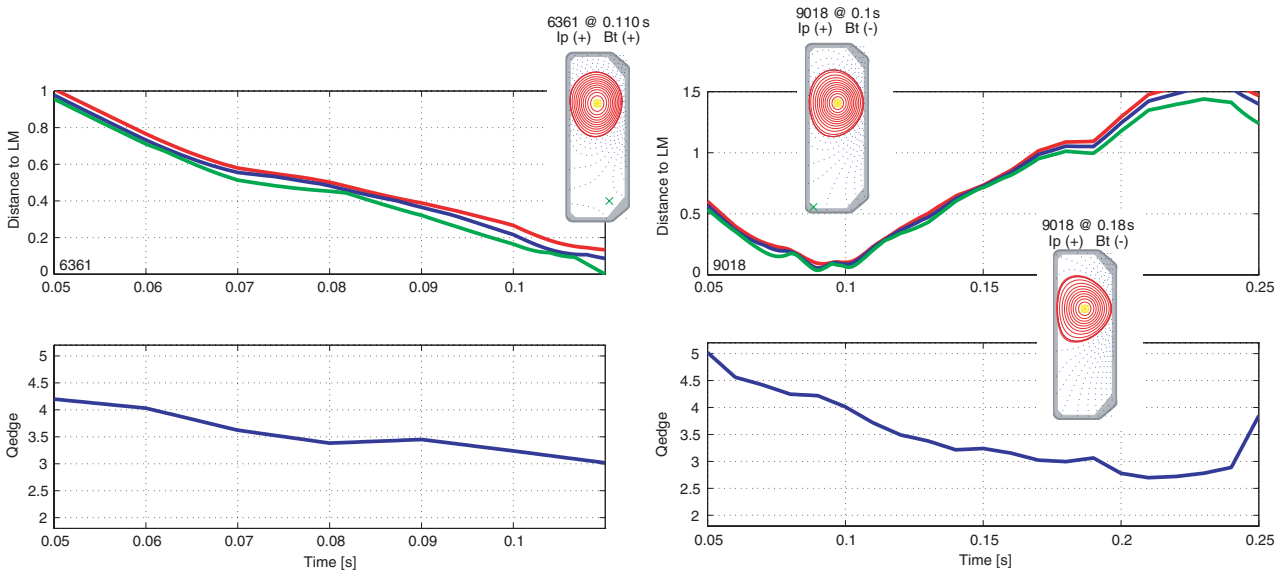

Figure 59. Time evolution of the distance to the l-closest locking modes (LM, see [395] for details), $l = 1,\ 3,\ 5$, shown together with the evolution of q-edge during pulse ramp up. In a standard case (on the left) modes lock at ∼0.11 s (as the distance to LM tends to zero). On the right, adjusted plasma shaping prevents locking and ensuing disruption [395].

always avoid disruption, since the typical plasma evolution time scales, such as the energy or particle confinement times, are significantly shorter than ramp down times. However, depending on the monitor limit and ramp-down parameter settings, ramp down can often reduce the severity of the disruption or VDE that ultimately occurs. In JET, for example, the soft-landing procedure can be initiated by any of several 'protection' monitor signals that are available. Monitored plasma characteristics may include the MHD activity level, radiated power fraction, x-ray radiation level and plasma density relative to the Greenwald limit.

Soft-stop actions can reliably effect machine protection and reduce disruption severity (see discussion of pre-disruption thermal energy loss in section 3.2). However, initiating soft-stop action can prematurely terminate pulses which are, in principle, 'repairable'. For example, many of the common conditions that lead to disruption—enhanced MHD activity and mode locking or strong divertor detachment—can be counteracted with intervention or 'repair' measures. One of the more ordinary reactions is to reduce, or stop, external gas or pellet injection if the plasma density becomes too high.

A more elaborate technique to forestall mode locking has been developed on TEXTOR [397]. A real-time cross-correlation method is utilized to detect growing $m = 2$, $n = 1$ modes. The ECE channels measure the electron temperatures close to the $q = 2$ surface, on both the high- and the low-magnetic field sides of the plasma column. If the cross-correlation coefficient exceeds a pre-programmed value, the injection of a tangential neutral beam is triggered, leading to an acceleration of the toroidal plasma rotation and reversal of mode growth (figure 60). This action not only prevents disruption but can also restore the enhanced plasma confinement regime existent before the onset of the mode growth.

Other examples of real-time repair actions include stabilization of (2,1) tearing modes by ECRH injection in JFT-2M [399] and RTP [400], control of mode locking by

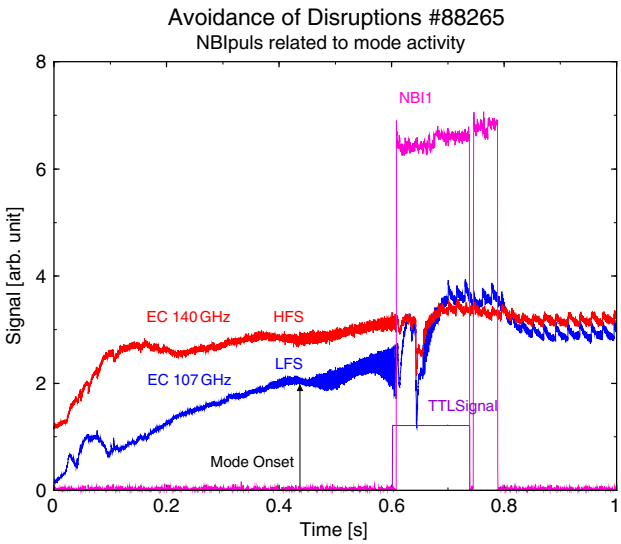

Figure 60. Time traces of high-field-side (HFS) and low-field-side (LFS) ECE channels. Detection of a high correlation coefficient initiates tangential NBI to speed up toroidal rotation and mode growth is reversed (TTL = transistor transistor logic signal derived from the ECE which triggers the NBI) (reprinted from [398]. Copyright 2001, with permission from Elsevier).

NBI-induced plasma rotation in JT-60U [401], control of the plasma stored energy in advanced tokamak plasmas in JT-60U [402] and DIII-D [254], control of resistive wall modes in DIII-D [190], heating power control of NTM onset or peaking of the pressure profile in JET [403], control of gas injection by monitoring the edge electron temperature via Langmuir probes in TORE SUPRA [404] and increase of NBI heating power and cessation of gas fuelling upon detection of 'deep divertor detachment' in ASDEX Upgrade [393]. In addition to these already implemented examples, emerging developments in 'active' MHD spectroscopy allow direct measurement of the damping rate of stable MHD modes [405, 406], and thus in

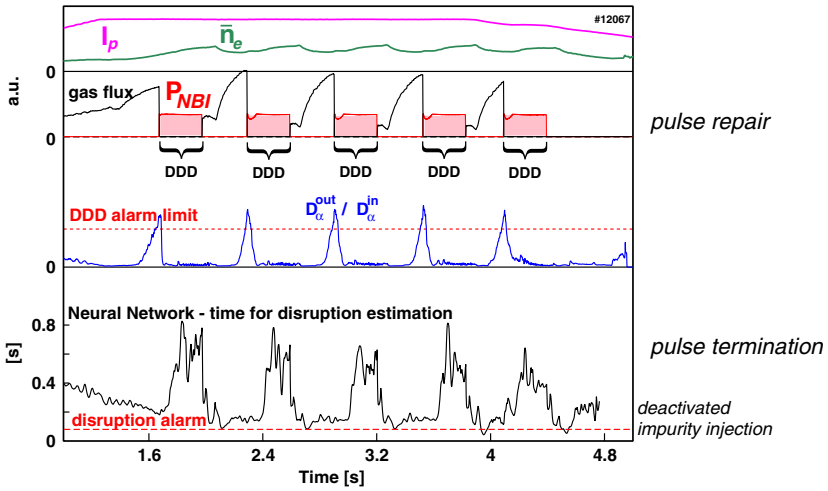

Figure 61. Competition of two independent machine protection processes (a pulse repair and a pulse termination reflex, see text). This example illustrates the need of a real-time supervisor to assure a clear control and protection hierarchy (reprinted from [408]. Copyright 2003, with permission from Elsevier).

principle provide a real-time method to measure the approach to a stability limit.

The various types of discharge repair actions noted are initialized via a discharge pulse control system and normally act independently from each other and/or the standard pulse programme control logic. This can lead to situations where two or more repair actions interfere. This is demonstrated in an example from ASDEX Upgrade in figure 61. The pronounced increase of the density in the Ohmic phase initiated the deep divertor detachment (DDD) repair action [393], which interrupts the standard pulse programme, increases the level of auxiliary heating power and closes all gas injection valves to restore stable plasma operation conditions. The DDD action is successful in repairing the discharge and preventing disruption. However, the estimated time for disruption, calculated in parallel by a real-time neural network, indicates a need to take mitigating action (neon impurity injection) to avoid excessive disruption-produced forces on the vacuum vessel. In the example shown, the injection-triggering output of the neural network was inhibited. If the output had not been disabled it would have triggered a (mitigated) disruption. This example shows clearly that a real-time supervisor is needed to ensure an unequivocal plasma control and machine protection hierarchy. This could be guaranteed, e.g. by a plasma regime identification algorithm [407], which dependently allocates a pre-defined protection action hierarchy on the basis of well-identified plasma and plasma operation states.

3.6.3. Disruption mitigation. The rationale and physics basis reasons for why disruption mitigation will be essential for the reliable operation of an ITER-class tokamak have already been extensively developed in the ITER Physics Basis [1] and in the preceding sections of this chapter. The fact that disruption mitigation also introduces sometimes-conflicting requirements on the method(s) to be deployed has also been developed in the IPB and in preceding sections. What follows below comprises a presentation of available methods and options for disruption mitigation in ITER. Some of these methods have been tested in present experiments; others constitute proposals for methods and/or technologies that remain to be realized and tested in a tokamak. As the conclusion of this section indicates, given the present uncertainties about how presently-tested and proposed methods will (can) work in ITER and given the already-noted fact that the conflicting requirements for disruption mitigation in ITER may force choices as to mitigation priorities, continuing pursuit of a range of options and technologies is justified.

Methods of disruption mitigation. In certain types of disruptions or loss-of-control VDEs, plasma control actions taken immediately upon onset may be able to at least partial mitigate some consequential effects, especially those associated with the current quench and vertical instability phases of such events. For example, experiments in JT-60U have demonstrated the possibility of mitigating vertical plasma motion in a major disruption by executing a rapid shift of the plasma vertical position immediately after the thermal quench is detected [334]. Experiments in several devices (e.g. [324, 334]) have shown that early injection of impurity species into a developing VDE reduces peak halo currents more than later injection. While the relatively long time scale (~0.5 s) for penetration of plasma equilibrium control fields through the ITER vacuum vessel may limit external control response to a VUD, the concept of using 'slow' impurity injection into a developing VDE to reduce halo currents appears to have merit. External control response to disruptions or VDEs that produce appreciable runaway current conversion is also critical to being able to sustain control of the resulting current channel until the runaway content can be benignly terminated (see section 3.4 and also chapter 8 of this issue [373]).

Mitigation of many of the most immediately-damaging disruption effects (including heat flux conducted along field lines, halo current loads and runaway electrons) requires rapid radiation of the plasma thermal and magnetic energies (as discussed below) coupled with the timely delivery of large numbers of electrons to the plasma volume during the current decay. As section 3.4 details, massive electron delivery immediately following the onset of the ITER current quench is likely critical to avoiding avalanche conversion of thermal

current to runaway current and the subsequent problems of then having to (very carefully) terminate the runaway current. Both these requirements—for radiating impurity delivery and for electron delivery—can be satisfied by injection of impurity atoms. Noble gas impurities are favoured owing to their otherwise benign interaction with in-vessel surfaces and ease of subsequent removal. The electron delivery and removal ease requirements can also be satisfied by injection of hydrogen or deuterium (although H_2 injection in ITER will introduce a non-reactive species that can compromise subsequent DT reactivity). Several methods of delivery are possible, all of which are conceptually capable, in an ITER embodiment, of delivering, within a few milliseconds, the large quantity ($\sim 10^{25}$ atoms or molecules) of impurities H_2 or D_2 needed.

Pellet injection. One obvious method for rapidly inserting impurities into a tokamak plasma is injection of impurity 'killer' pellets. Injected species used in killer pellet (KP) tests in present experiments have included cryogenic H_2, D_2, Ne, Ar, Kr-doped cryogenic D_2 and Ag-doped polyethylene. The efficacy of the KP approach in reducing divertor energy deposition and halo current magnitude and toroidal asymmetry has been demonstrated in many devices. Reductions of 25–95% in thermal flux conducted to the divertor, 50–75% in peak halo current magnitudes and reduction in halo current toroidal peaking factors (TPFs) to unity (e.g. [332, 365, 409, 410]) have been demonstrated. Modelling has also contributed to improved understanding of disruption mitigation by KP injection [364]. Various possible pure and mixed-species pellets have been investigated in these modelling studies, including D_2, He, CH_4, Ag-doped H, Kr, Xe, Ne and Ar. Injection of low-Z materials such as D_2 and He is found to be less effective in mitigation than the high-Z radiators such as Ne and Ar (e.g. [366]).

Killer pellet injection velocities are typically $\leqslant 1\,\mathrm{km\,s^{-1}}$. As with standard plasma fuelling by injection of D_2 pellets, the penetration depth for KPs based on present-day single-stage gas gun or centrifuge technologies will be much smaller than the ITER plasma minor radius. Both low-field-side (LFS) and high-field-side (HFS) injections are possible. Higher velocities and deeper initial penetration are attainable with LFS injection, but polarization drift (see e.g. [411]) tends to inhibit further inward transport of the injected impurities. The velocity limitations imposed by HFS guide tubes limit usable HFS velocity, but penetration, even with reduced velocity is deeper owing to polarization drift and resulting inward transport of the injected impurities. Anomalous pellet material transport to the plasma centre has also been observed with LFS injection [412].

The principle shortcoming of cryogenic pellet injection is the frequently-observed production of runaway electrons (figure 62), which arises owing to the high local electric field that pellet injection and subsequent ablation produces [365,413]. While the details of how the runaways are produced involves careful consideration of the local impurity deposition effect and modification of the electron distribution function, the inherent propensity of KP injection to produce runaways can be understood from the simple fact that the number of impurity electrons added by typical cryogenic pellets (e.g. in DIII-D, $\sim$1–2 mm diameter Ar pellets containing $\sim 5 \times 10^{20}$ atoms) is insufficient to maintain $E \leqslant E_c$, where $E_c = (4\pi e^3 n_e \ln\Lambda / mc^2)$ is the critical electrical field that governs whether or not the avalanche grows (see section 3.4 and the massive gas injection discussion that follows below).

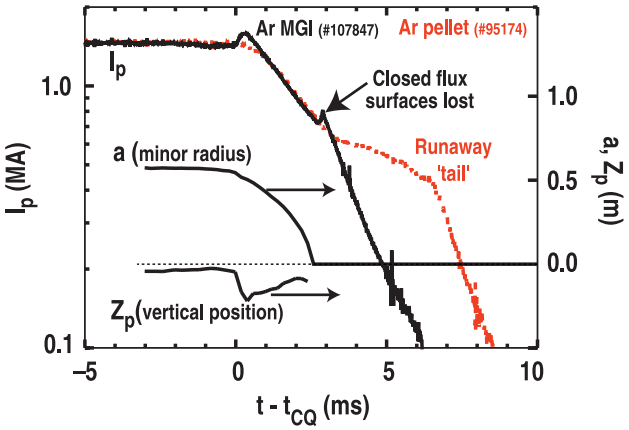

Figure 62. Plasma current, size and position waveforms for Ar killer pellet and Ar massive gas injection in DIII-D. A $\sim$0.6 MA runaway current 'tail' develops following the Ar KP injection (dotted curve). Adapted from data originally presented in [365] and in [366].

Massive gas injection. Experiments conducted since the writing of the 1999 IPB [1] have shown that injection of massive amounts of impurity and/or H_2 or D_2 neutral gas into tokamak plasmas can provide the beneficial disruption mitigation effects of impurity killer pellets but with a reduced or negligible tendency to produce runaways. Here the terminology massive gas injection (MGI) is used to denote disruption mitigation methods that employ a gas injection system (typically a small, high-pressure gas reservoir coupled to a fast-opening valve) to produce a subsonic or trans-sonic gas jet that is capable of delivering a large quantity of neutral gas, typically more than 10^{22} atoms ($\sim$0.1 atm litre $= 100\,\mathrm{Pa\,m^3}$) to the torus vacuum vessel within 10 ms. Table 11 shows the injection parameters and species used in the various DIII-D MGI studies ([324, 365, 366, 371]) conducted during the period 1997–2005. Injection rates for ASDEX Upgrade and JT-60U experiments are $\sim$3 and $\sim$30 times lower, respectively. Automatic initiation of neon MGI taken up on warning of impending disruption is now routinely used to limit vertical force on the vacuum vessel in ASDEX Upgrade [414] (see section 3.6.2 above).

Experiments with noble gas MGI on many devices (e.g. [362, 365–367, 371, 414, 415]) have demonstrated the highest level of mitigation of thermal and electromagnetic loads observed in tokamaks, both in major disruptions and VDEs. Measurements in DIII-D show that the propagation of the cooling front to the plasma centre occurs with a radial propagation velocity that is commensurate with the room-temperature sound speed of the injected species (He, Ne or Ar) and bolometric data show essentially complete radiation of the plasma stored energy [324]. In ASDEX Upgrade and DIII-D, halo current magnitudes are typically reduced by more than 50%, and toroidal peaking factors are reduced to unity, resulting in a typical reduction in electromagnetic forces of more than 75% [331, 414]. The reduction in halo current magnitude and TPF obtained in MGI experiments is consistent with the fast current quench, reduced vertical instability and

Table 11. Gas injection capabilities for DIII-D MGI experiments (1997-2005)[a].

Gas	Flow (10^5 Torr L s^{-1})	Flow (10^4 Pa m^3 s^{-1})	Total injected quantity (10 ms) (10^{22} atoms or molecules)
96% H_2 + 4% Ar	~ 2.5	~ 3.3	~ 8.4
D_2	3.7	4.9	12.4
He	3.5	4.6	11.6
Ne	1.3	1.7	4.59
Ar	0.93	1.22	3.28

[a] Experimentally measured equilibrium flow rates, typically obtained for 70 atm reservoir pressure. Some variation in valve and/or injection configuration has occurred over the data collection period.

Table 12. KPRAD model calculations and experiment data for DIII-D disruption mitigation experiments (data adapted from [366]). Here n_{Ar} is the injected argon density, $\langle Z \rangle$ is the average charge state, T_e is electron temperature and τ_{CQ} the L/R current decay time.

	Model					Experiment	
Mitigation means	n_{Ar} (m^{-3})	$\langle Z \rangle$	T_e (eV)	E/E_c	τ_{CQ} (ms)	$\langle Z \rangle$	τ_{CQ} (ms)
Ar pellet	2×10^{19}	5.9	7.5	312	2.0	5.7 ± 1	2.1 ± 0.1
Ar MGI	2×10^{21}	0.45	1.46	3.8	1.38	0.3 ± 0.1	1.8 ± 0.1

cooling/symmetrization of the ex-plasma halo region that is also obtained in naturally occurring fast-current-quench 'radial disruptions' (see section 3.3.2).

The high electron densities that MGI is capable of producing can approach the magnitudes necessary for collisional suppression of runaway avalanching (see section 3.4). Table 12 compares the impurity densities and impurity ionization states (volume-average plasma Z_{eff}) estimated by Whyte *et al* [366] to have been obtained for otherwise comparable DIII-D argon killer pellet and MGI examples (e.g. the cases illustrated in figure 62). The calculations were effected with a self-consistent 0-D KPRAD impurity radiation model. The estimated and measured plasma Z_{eff} and current quench times (exponential decay time constant, τ) are found to be in good agreement, lending support to the hypothesis that the temperature and ionization state of the after-injection killer pellet and MGI plasmas are determined, as they often are for natural disruptions, by radiation-loss versus Ohmic-input power balance considerations. Owing to their high impurity density, MGI plasmas attain a low-T_e equilibrium state, where the Ar impurity is only partial ionized and where all the injected electrons are presumably present (mainly in bound form) to contribute to collisional suppression of runaway growth.

Further DIII-D MGI experiments conducted with a modified injection system designed to provide higher jet pressure at the plasma surface [371] have shown that the direct penetration of injected Ar gas into stationary H-mode plasmas is limited to a few cm. Fast-gated camera images in Ar-I light indicate that the neutral deposition is initially localized near the injection location and that there is no indication of singly ionized Ar penetration far into the pedestal edge gradient region. A similar lack of direct penetration is observed for low-temperature Ohmic plasmas and during the after-thermal-quench phase of MGI (where the estimated plasma temperature has dropped to ~2 eV). Nevertheless, despite this lack of observable direct neutral penetration, fast bolometry indicates that high (~90%) fractions of volumetric thermal

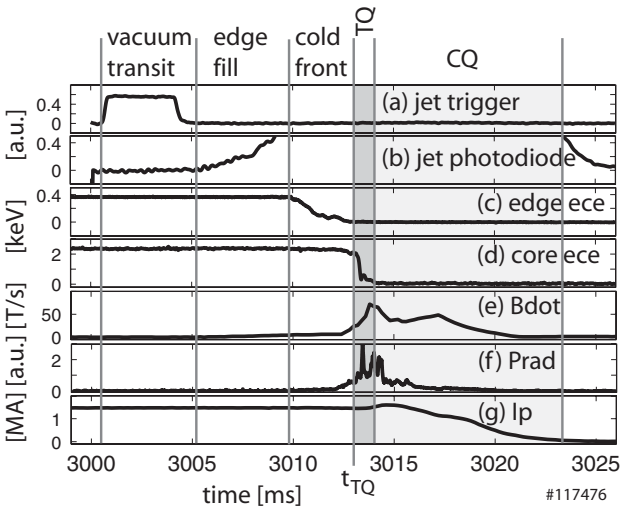

Figure 63. Fast plasma shutdown sequence in DIII-D (Ar MGI) [371] showing (*a*) the gas jet solenoid current, (*b*) visible emission, (*c*) second harmonic edge ECE, (*d*) core ECE, (*e*) magnetic fluctuations, (*f*) radiated power and (*g*) plasma current.

energy radiation are achieved within ~2 ms of initiation of significant impurity delivery; this efficient radiation of thermal energy appears to be facilitated through a combination of rapid impurity-ion mixing and fast heat transport, both driven by the onset of large-scale MHD activity that follows the initial arrival of gas at the plasma surface. Figure 63 shows the sequence of events that is observed in all DIII-D MGI examples. Initial arrival of gas at the plasma edge (precursor flow $\ll$ equilibrium flow [416]) is followed by the onset of edge cooling and a slight drop in plasma current (indicative of a narrowing of the current profile), abrupt increase in MHD activity (2/1 and other $n = 1$ modes), enhanced volumetric radiation and onset of thermal quench and a disruption-like positive plasma current spike that follows the thermal quench onset. A rapid current quench that is usually runaway-free follows. A non-

Table 13. Parameters for a single-species ITER MGI system.

Value or requirement	Units	D_2	Ar	Basis or assumption
Sound speed (v_s)	$m\,s^{-1}$	935	320	300 K
Time to reach plasma surface	ms	3.1 (5.3)	9.4 (15.6)	3 m (5 m) to plasma; propagation at v_s
Additional delay	ms	2	2	Valve trigger and opening delay, etc
Minimum look-ahead time for pre-emptive action	ms	5.1 (7.3)	11.4 (17.6)	Sum of propagation and delay times, assumes deployment before natural TQ onset is required
N_{RB}	Atoms or molecules	1.75×10^{25}	1.94×10^{24}	To achieve $n_e = n_{RB}$, for 830 m^3 plasma volume, 100% assimilation
Delivery time (t_{del})	ms	9	9	$4^* t_{del} = t_{CQ} = 36$ ms; $t_{CQ}/S^* = 1.7$ ms m^{-2}
Required average flow rate	s^{-1}	1.94×10^{27}	2.16×10^{26}	Assumes prompt rise
Required average flow rate	Pa m^3 s^{-1}	7.2×10^6	8.0×10^5	Assumes prompt rise
Required average flow rate	Torr L s^{-1}	5.5×10^7	6.1×10^5	Assumes prompt rise

ideal MHD numerical simulation (NIMROD) has been used to recreate similar experimental features of MGI mitigation on Alcator C-Mod [417]. Forced plasma cooling at the plasma edge was found to trigger robust MHD modes, which then led to development of stochasticity and the loss of internal plasma energy by heat conduction into the cold boundary plasma, without any impurity particle penetration deep into the core plasma. This result verifies the important coupled role of plasma cooling and MHD, in disruption mitigation.

Modifications to the DIII-D injection system made over the course of the latest set of experiments that changed the rate of rise of initial gas delivery were found to shorten the overall time scale of the sequence of events that leads up to thermal quench onset, but were found to otherwise have little effect on the subsequent current decay phase. The sequence of events and characteristics observed following MGI in DIII-D closely resemble those observed for natural 'high-density' fast disruptions. However, the MGI-initiated 'disruption' exhibits (more) benign divertor energy loading and halo current characteristics and favourable runaway electron mitigation characteristics relative to otherwise comparable killer pellet initiated fast shutdown (note that the KP shutdown shown in figure 62 does not show a current spike or other evidence of internal reconnection).

A detailed physics understanding of the mechanisms for gas jet penetration and ion mixing in the thermal quench and current quench phases of MGI is not yet available but is a topic of present research. Gas injection also offers a possible means for raising impurity densities during the CQ itself. Other techniques, e.g. high-velocity pellets or liquid jets, are predicted to simply pass through the cold after-TQ plasma without appreciably further raising the density or slowing runaway electrons. As the studies of methods for ITER indicate, meeting the competing requirements of mitigation for the various potentially damaging disruption effects may require combined or separate injection of different species of impurities at different points in the shutdown process.

Liquid jets. It has been proposed that one method of achieving sufficient impurity penetration for disruption mitigation in a reactor-grade plasma is to use cryogenically cooled liquid jets. Calculations suggest that this method may offer the capability of depositing sufficiently large quantities of impurities well within the plasma [418]. Penetration may actually be so efficient as to require careful control of jet characteristics to prevent 'shine through' and damage to PFCs from direct jet impact. Although the full technology for cryogenic liquid jet injection has not been demonstrated, prototypes of room-temperature liquid (water) injectors with relevant pressures and operating parameters have been produced and demonstrated (e.g. [419]).

Other concepts. Mitigation methods based upon multiple HFS-launch pellet injection (a sequential burst of moderate-sized pellets, intended to facilitate incremental penetration, [364]) and hypervelocity injection of dense compact plasma toroids or spheromaks [420] have been proposed for ITER and reactor application. Both methods offer a potential solution to the problem of effecting direct and immediate delivery of radiating impurities to the plasma core. Testing of either concept in present tokamaks and development of ITER-applicable injection technologies remain as future research tasks.

Application to ITER. There are multiple challenges in developing reliable disruption mitigation methods for ITER. Successful scenarios must bring together real-time stability analysis and timely disruption prediction, reliable and timely triggering of the mitigation system and mitigation means that provide unequivocal runaway electron (RE) suppression. Above all, the underlying principle must be that the mitigation scheme 'does less harm', both with regard to recovery of plasma operations and maintaining the viability of plasma-facing components than allowing unmitigated occurrence of disruption. Experimental and modelling advances in detection (3.6.1), passive avoidance and precision/intelligent control (section 3.6.2) and mitigation means that can act to inhibit runaway conversion point the way to an ITER disruption mitigation scheme based on massive impurity and/or hydrogenic injection. The single-species MGI approach described above, which so far has been successful in mitigating most of the consequences of disruptions in present tokamaks, offers promise as an ITER disruption mitigation means. Table 13 gives a set of elementary 'design basis requirements' for a single-species MGI system for ITER. Possibilities for either low-Z (D_2) and moderate-Z (Ar) injections are shown; requirements for intermediate-Z species (or mixtures of higher-Z noble gases with D_2) lie between these extremes.

Moderate-Z (Ar) injection reduces the requirement for injection quantity and flow rate because more electrons arrive

for each gas atom [366] but requires a longer look-ahead time for pre-emptive deployment. Low-Z (D_2) injection increases the gas quantity and flow requirements but appreciably reduces the look-ahead time requirements. The estimates in table 13 assume that the MGI system valve (or valves) will, for reasons of nuclear radiation resistance and service access, need to be located between 3 and 5 m from the plasma surface.

There are also further issues in choosing between moderate-Z and low-Z injection: the gas delivery rate for a given valve and 'jet tube' design (see e.g. [371]) scales as v_s (the gas velocity), while the gas flow rise time (the time for the gas flow through the 'jet tube' to reach full equilibrium value) scales as v_s^{-1}. Hence the system flow capability for D_2 is ~3 times higher than the corresponding Ar capability and rise time will be ~3 times shorter. To first approximation, these gas delivery factors offset the 9 times increase in electron delivery rate that Ar affords relative to D_2. Also, given that the measured rise-time for Ar for the 1.3 m long DIII-D MGI system used for the reference [371] experiments is ~10 ms [416], rise-time limitations imposed by a 3–5 m long ITER system may adversely impact the prompt electron delivery efficacy for the moderate-Z or high-Z injection.

The parameters in table 13 are based on a requirement to achieve the Rosenbluth no-avalanche density within $t_{CQ}/4$ and incorporate no explicit assessment of the radiation cooling effectiveness of the injected impurities in mitigating the divertor thermal energy loading or in achieving sufficient after-TQ radiation to obtain a prompt current shutdown (with low halo current). While the $t_{CQ}/4$ specification is somewhat arbitrary, its intent is to ensure that no significant runaway electron avalanching occurs in the early phases of the current quench, when the driving electric field is highest. And while there is no significant question about the ability of the quantities of injected Ar (or Ne) specified in table 13 to efficiently radiate the full 350 MJ of ITER thermal energy within a few ms, there are concerns, first identified in the IPB, that too-effective radiation during the initial edge-cooling onset phase of MGI (before internal reconnection and TQ develop) can lead to excessive local or global energy loading on ITER first-wall surfaces.

Here elementary calculations show that the minimum deposition time for uniform radiation of the full W_{th} to a beryllium first-wall should be $\geqslant$~0.6 ms if surface melting is to be avoided. Local peaking factors or pre-heating of the first-wall surface will increase this minimum requirement. While some tailoring of the initial radiation cooling attributes of a D_2-dominated system may be possible by judicious admixture of higher-Z noble gases [367], present physics understanding of the combination of gas jet rise-time characteristics, neutral gas penetration, impurity ionization and transport along flux surfaces and overall radiation dynamics is insufficient to make detailed estimates of exactly what will happen to the first-wall surface from ITER MGI deployment. Present physics understanding is also insufficient to determine whether injection of the quantities of gas specified in table 13 will, in fact, result in effective suppression of runaway avalanching in the plasma core during the ensuing current decay.

Given these uncertainties about application of MGI for thermal energy and runaway mitigation, the choice of the optimal species and delivery technique(s) for impurity injection disruption mitigation in ITER remains open: as the previous discussions have shown, high-pressure gas jets (MGI), liquid jets and single and multiple (sequentially injected) cryogenic pellets all remain as possible options. At a minimum, any option or combination of options selected must satisfy a fundamental requirement of being able to deliver the impurities required within a time shorter than the delay between detection of impending disruption and onset of a natural disruption TQ. Since most of the proposed delivery systems are limited to helium-sound-speed (10^3 m s^{-1}) injection velocities, if the impurity injector and associated hardware are to be located outside the ITER toroidal field coil (on an equatorial port), the minimum transit time to $r/a = 0.5$ in the plasma will be $\geqslant$6 ms. The combination of this minimum propagation time with a 2–4 ms action time (opening of a valve, etc) results in the need for an 8–10 ms look-ahead time, plus prompt action of the disruption detection (prediction) and control authorization software. While the detection of many types of ITER disruptions more than 10 ms before occurrence appears feasible (sections 3.6.1 and 3.6.2), detection of impending fast β-limit disruptions within 6 ms may (owing to diagnostic limitations or fast MHD growth rates) be less feasible.

In this case, to minimize delay time, the injector hardware must be placed as close as possible to the plasma surface (so that propagation times approach 1 ms), mechanical and electronic system activation delays must be minimized and high velocity injection schemes are favoured. Considerations of simplicity, volume and reliability argue for a small-volume high-pressure gas reservoir and a fast-acting valve with a minimal nozzle or jet tube. Compact and/or magnetic-field-tolerant systems of this nature have already been implemented in ASDEX Upgrade [386] and TEXTOR [415]. However technology development is required to design an ITER-qualified system that is also compatible with both the high magnetic-field and nuclear radiation and tritium environments and is otherwise capable of meeting the short activation time, fast rise time and total gas quantity requirements.

Recent observations of MGI gas jet penetration in DIII-D [371] and models [421], developed to explain the lack of direct penetration even for low-field, low T_e plasmas, indicate that the local jet ram pressure (ρv^2) at the plasma surface must approach the local magnetic field pressure ($\sim B^2/2\mu_0$) for effective penetration. If this requirement applies for ITER, then gas jet surface pressures of 10 MPa (100 bar) will be required for direct penetration. It is unlikely that a compact/fast-acting gas jet system capable of generating this pressure at the plasma is possible. MGI in ITER will therefore need to depend on surface delivery, with subsequent MHD mixing of the surface-deposited impurities (as appears to be the case in present experiments).

Pressures of ~100 bar are more easily attained with liquid jets and calculations have shown that liquid He and/or liquid H_2/D_2 jets can penetrate into the ITER core [413, 418]. While further assessments of the technological feasibility of cryogenic liquid jets are needed and while concerns remain about jet fragmentation (owing to hydrodynamic stability), once the jet stream enters the plasma, liquid jets continue to offer a potentially attractive means for rapidly delivering impurities to near the centre of the ITER plasma.

Finally, the calculations of Kuteev *et al* [422] have shown the conceptual (ablation physics) feasibility of achieving deep penetration with the LFS launch of high-velocity (3–5 km s^{-1}) cryogenic noble gas pellets with diameters ~10 mm. Velocities of this magnitude can be achieved with two-stage gas-gun launchers, albeit with a need for appreciably more look-ahead time than is required for a single-stage pellet gun or gas jet and a need to reinitialize the launcher system after each use. Again, further technology assessments are needed and there is no present experience with the interaction of such large-scale hypervelocity pellets with a correspondingly large-scale, high-temperature and high-density tokamak plasma. The considerations noted above as to the possible need to 'tailor' the pellet species and spatial and/or time-profile of the delivered impurity (or impurities) to obtain a first-wall acceptable radiation history also apply.

Taking the above considerations together, there is some basis for cautious optimism about the development of disruption mitigation schemes that will be applicable to ITER. The need for effective mitigation, including mitigation of runaway electron conversion, is unequivocal, and runaway mitigation will be needed once high-current plasma commissioning operations commence in the hydrogen phase. With regard to urgent research needs in present experiments, the most pressing areas for research appear to be related to detection, gas jet penetration and subsequent impurity mixing throughout the plasma volume and strategies to avoid or minimize radiative melting of the first-wall. While gas injection systems appear to offer many advantages for eventual utilization in ITER as primary and back-up means for disruption and/or runaway electron conversion avoidance, concerns about the ability of gas jets to effect direct control over impurity delivery and in-plasma distribution during the onset phase of mitigation suggest that alternative liquid jet, multi-pellet and hypervelocity pellet injection approaches should also be pursued.

3.7. Summary and R&D needs

There has been substantial new data and progress in understanding the physics basis for disruptions and their consequential effects in an ITER-class tokamak. In addition, the size and plasma energy content of ITER have been reduced (relative to the ITER-EDA design), resulting in some relaxation of the divertor PFC surface energy loadings. There has also been substantial progress in developing disruption consequence mitigation schemes and in improving prediction and avoidance strategies to minimize the number of disruptions that will occur during ITER operation.

The overall understanding—summarized in the ITER Physics Basis [1]—of the MHD instability processes and causative factors that contribute to onset of disruption remain largely unchanged. Conventional pictures of the internal triggering mechanism for major disruption, based on magnetic reconnection and non-linear growth of helical instabilities (cold bubble ingress) in the low-shear central region, are now being supplemented by internal ballooning/interchange models that better explain the observed insensitivity of internal thermal equilibriation times to initial plasma temperature. It is likely that both mechanisms play a role in disruption development. The emergence of full 3D MHD dynamic equilibrium modelling capabilities offers prospects for quantitative predictions of ITER disruption scenarios and consequences. The application of these models to interpretation of data in present tokamaks will likely aid in sorting out some of the presently obscure internal cause-and-effect dynamics of the precursor-growth and thermal-quench-onset phases of disruption. Three-dimensional MHD simulations will perhaps also aid in clarifying the cause(s) for the wide range of data scatter in the existing current quench and halo current databases.

Data quantifying the range of physics and hardware/operations 'causes' for disruptions demonstrate that all disruptions have identifiable physics and/or operations causes and that there are often clear distinctions in the precursor and development phases of disruptions initiated by a plasma edge energy balance deficit ('cold-edge' or density-limit-onset disruption) versus those initiated by rapidly growing internal ideal-MHD instabilities (beta-limit or ITB pressure-gradient disruptions). There are also 'thermal-collapse' disruptions, initiated by the onset of slowly-growing resistive MHD instabilities (e.g. NTMs), that share many similarities with cold-edge disruptions. Cold-edge and thermal-collapse disruptions have long development phases, with ample before-onset diagnostic indication(s) of impending disruption; ideal-instability-onset disruptions have much shorter development phases and only limited external indications that disruption is imminent. The difference in onset characteristics impacts the types and reliability of predictive methods that can be used to take timely before-disruption avoidance or mitigation action. Both single-parameter deterministic indicators and multi-parameter neural net methods have been successful in reliably predicting many, but not all, types of disruption onset in present tokamaks, and both methods hold promise for ITER. Concerns do exist, however, about the *a priori* need for data from disruptive ITER plasmas to 'train' an ITER neural network, though there has been some initial progress on this issue. For both methods, achieving an acceptable balance between prediction reliability and 'false-alarm' rates is also an issue.

Statistical studies of the physics and hardware causes and 'types' of disruption continue to support the understanding documented in the IPB that the overall frequency of disruptions in tokamaks conducting a wide-ranging programme of plasma development and exploration studies is about 10%. Disruptivity (frequency of disruption per pulse) tends to be higher during exploratory phases or campaigns to develop higher-performance plasmas, but higher rates of disruption are not always correlated only with proximity to one or more of the three traditional operations limits. There is also emerging statistical evidence that once the prescription and control means for achieving a stationary high-performance plasma state is obtained, there is no further tendency for such plasmas to disrupt as the duration of the stationary phase is extended. The reproducibility of the plasma 'start-up phase' (which here denotes the entire approach to stationary conditions) can, with adequate pre-discharge wall conditioning and hardware/control reliability, approach 100%. These findings support the expectation that a combination of 'passive' and 'precision control' disruption avoidance strategies coupled with well-qualified prediction capabilities will be successful in

contributing to the achievement of low disruptivity in ITER, even in exploratory regimes.

Expectations for avoiding disruptions and/or being able to successfully mitigate many of their effects notwithstanding, present data continues to support the need to design the ITER vacuum vessel and in-vessel component structures to withstand the electromagnetic loading effects of a credible number of unmitigated 'limit-case' disruptions and VDEs. Newly-acquired and re-interpreted data on the current decay time scaling indicates that the lower bound on area-normalized current decay time is approximately 1.7 ms m^{-2}. New data on halo current magnitude and toroidal asymmetry (TPF) confirm the continued use of the IPB-identified guidelines of $I_h/I_{p0} \leqslant 0.4$ and I_h/I_{p0}*TPF $\leqslant 0.75$, albeit with increasing evidence that TPF > 2 can be excluded as a design basis. Emerging 3D MHD code modelling capabilities also hold the promise of being able to make first-principles-based simulations of halo current and 3D MHD effects in ITER disruptions and VDEs.

New and more comprehensive data on the internal redistribution and subsequent loss of plasma thermal energy (W_{th}) to divertor and other in-vessel PFC surfaces—in combination with the twofold reduction in the ITER thermal energy loading ratio—has significantly improved prospects for avoiding disruption-produced thermal erosion of divertor PFC surfaces. Thermal energy loss and deposition observations in present tokamaks show that while internal W_{th} redistribution times for ITER are still projected to be ~1 ms, the timescale for divertor surface thermal energy deposit will be significantly longer, perhaps as long as 10 ms. Significant broadening of the divertor energy deposition area, by a factor of about 10, is also now expected, and there is increasing present-experiment evidence for 'pre-divertor' convective and/or radiative deposition of a significant fraction of the plasma thermal energy on the divertor entrance and first-wall surface areas. Finally, accumulating evidence in present tokamaks suggests that significant before-disruption thermal energy loss may be expected in at least some classes of ITER disruptions (e.g. those with slowly developing initial phases). While there is presently insufficient understanding to quantify how these various W_{th}-mitigation factors will combine in ITER, the prospects of obtaining reduced or zero divertor PFC erosion during many ITER disruptions is now significantly enhanced relative to previous prospects arising from the 1999 ITER Physics Basis [1].

Runaway electron conversion continues to be serious concern for ITER. Already, in JET, half of the thermal pre-disruption current can be converted into runaway current, with much of the runaway population generated by the Coulomb avalanche multiplication process. Avalanche gain in ITER will be much higher and simulations show that, without additional losses, a substantial fraction (70–80%) of the thermal current will be converted into runaway current following disruption or onset of the current quench phase of a VDE. Other simulations of the effect of the runaway energy deposition expected on the ITER first-wall or divertor surfaces indicate that an uncontrolled interaction of this magnitude of runaway current has the potential to produce PFC surface and substrate melting and erosion.

Two mechanisms to forestall runaway conversion in ITER are identified: (i) enhanced prompt runaway loss produced by large-amplitude magnetic fluctuations and (ii) increase of the plasma electron density (free + bound) to the Rosenbluth no-avalanche density, ~10^{22} m^{-3}, within a few ms after the onset of current quench. Both of these mechanisms present challenges to implement in ITER. For fluctuations, the challenges lie in how to rapidly apply externally generated helical fields or in how to induce natural (self-generated) MHD fluctuations, that are capable of effecting the prompt (~10^3 s^{-1}) loss rates needed to offset avalanche gain. For electron density increase, the challenges lie in being able to reach the Rosenbluth density before appreciable current decay develops: this will require the delivery of ~3×10^{25} electrons (= 1.5×10^{25} H_2 or 1.8×10^{24} Ar) to the ITER plasma in $\leqslant$10 ms. Options for achieving these high delivery rates include sequential multiple pellet injection, high-velocity liquid jets and massive gas injection (MGI). Of these three options, only MGI has so far been tested in present tokamaks at ITER-equivalent injection rates. While the ability of MGI (and also small 'killer pellets') to successfully reduce divertor energy deposition and halo current magnitude and asymmetry in present-tokamak testing is unequivocal, improvement of present injection system flow rate and rise time capabilities and tests in a large, high-current tokamak are needed before a definitive demonstration of reaching the Rosenbluth density for runaway prevention within a fraction of the plasma current decay can be demonstrated. Physics and extrapolation basis questions also apply to how the presently observed role of self-generated MHD fluctuations in promoting rapid 'MHD mixing' of the edge-ionized atoms that MGI produces will extrapolate to ITER and also, how, in situations where avalanche gain is high, self-generated fluctuations affect runaway multiplication.

Finally, there has been continuing progress in improving the scope and application of integrated disruption models. Two-dimensional dynamic equilibrium codes with various degrees of auxiliary models specialized for disruption and VDE modelling are now available and have been used to make self-consistent estimates of ITER electromagnetic loadings. However, further experimental validation, better self-consistent integration of the auxiliary models, expansion of dynamic modelling to include full 3D plasmas and 3D vessel and in-vessel component representations are needed to establish comprehensive disruption scenarios and evaluations of disruption mitigation in ITER.

In summary there has been steady progress in understanding disruptions and their consequential effects. Of particular note is progress in developing massive gas injection techniques for ameliorating disruption consequences. There does remain a requirement to develop better disruption prediction techniques, to better understand energy deposition on the first wall and divertor, to develop a truly predictive modelling capability for ITER and to fully understand mechanisms of how injected gas penetrates into the plasma and the capacity of such techniques to inhibit runaway electron formation.

References

[1] ITER Physics Basis 1999 *Nucl. Fusion* **39** 2137
[2] Gormezano C. *et al* 2007 Progress in the ITER Physics Basis *Nucl. Fusion* **47** S285–S336

[3] Loarte A. *et al* 2007 Progress in the ITER Physics Basis *Nucl. Fusion* **47** S203–S263
[4] Sauter O. *et al* 2002 *Phys. Rev. Lett.* **88** 105001-1
[5] Nave M.F.F. *et al* 1995 *Nucl. Fusion* **35** 409
[6] Porcelli F., Boucher D. and Rosenbluth M. 1996 *Plasma Phys. Control. Fusion* **38** 2163
[7] Kadomtsev B.B. 1975 *Fiz. Plazmy* **1** 710
Kadomtsev B.B. 1976 *Sov. J. Plasma Phys.* **1** 389
[8] Ottaviani M., Porcelli F. and Grasso D. 2004 *Phys. Rev. Lett.* **93** 075001
[9] Gimblett C.G. and Hastie R.J. 1994 *Plasma Phys. Control. Fusion* **36** 1439
[10] Sauter O. *et al* 1999 *Proc. Joint Varenna-Lausanne Int. Workshop on Theory of Fusion Plasmas, 18th Int. School of Plasma Physics (Varenna, 1998)* ed J.W. Connor *et al* ISSP-18 (Bologna: Editrice Compositori) p 403
[11] Lazzaro E. *et al Proc. 26th EPS Conf. on Controlled Fusion and Plasma Physics (Maastricht, The Netherlands, 1999)* vol 23J (ECA) p 381
[12] Angioni C. *et al* 2000 *Theory of Fusion Plasmas: Proc. Joint Varenna–Lausanne Int. Workshop, 19th Int. School of Plasma Physics (Varenna, 2000)* (Bologna: Editrice Compositori) pp 73–86
[13] Mantsinen M.J. *et al* 2002 *Plasma Phys. Control. Fusion* **44** 1521–42
[14] Angioni C. *et al* 2003 *Nucl. Fusion* **43** 455
[15] Angioni C. *et al* 2002 *Plasma Phys. Control. Fusion* **44** 205
[16] Angioni C. *et al* 2002 *Proc. 29th Conf. on Plasma Physics. and Controlled Fusion (Montreux, Switzerland, 2002)* vol 26B (ECA) P-1.118
[17] Goodman T.P. *et al* 2003 *Nucl Fusion* **43** 1619
[18] Goodman T.P. 1999 *Proc. 26th EPS Conf. in Controlled Fusion and Plasma Physics (Maastricht, The Netherlands 1999)* vol 23 (ECA) p 1101 P3.040
[19] Henderson M.A. *et al* 2001 *Fusion Eng. Des.* **53** 241
[20] Mück A. *et al* 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St. Petersburg, Russia, 2003)* vol 27A (ECA) P-1.131 and http://eps2003.ioffe.ru/
[21] Zohm H. *et al* 2003 *Nucl. Fusion* **43** 1570
[22] Zohm H. *et al* 2006 *Proc. 21st Int. Conf. on Fusion Energy Research 2006 (Chengdu, China)* (Vienna: IAEA) CD-ROM EX/4-1Rb
[23] Isayama A. *et al* 2002 *J. Plasma Fusion Res. Series* **5** 324
[24] Pinsker R.I. *et al* 2003 *Bull. Am. Phys. Soc* paper GO1.004 http://flux.aps.org/meetings/YR03/DPP03/baps/abs/S930004.html
[25] Hanada K. *et al* 1992 *Phys. Fluids* B **4** 2675
[26] Westerhof E. *et al* 2003 *Nucl. Fusion* **43** 1371
[27] Ikeda Y. *et al* 2002 *Nucl. Fusion* **42** 375
[28] Mück A. *et al* 2002 *Proc. 29th EPS Conf. on Controlled Fusion and Plasma Physics (Montreux, Switzerland, 2002)* vol 26B (ECA) P-1.037
[29] Pietrzyk Z.A. *et al* 1999 *Nucl. Fusion* **39** 587
[30] Isayama A. *et al* 2001 *Fusion Eng. Des.* **53** 129
[31] Petty C.C. *et al* 2003 *Nucl. Fusion* **43** 700
[32] Pochelon A. *et al* 2001 *Proc. 28th EPS Conf. on Controlled Fusion and Plasma Physics (Funchal, 2001)* (ECA) vol 25A p 1805
[33] Eriksson L.-G. *et al* 2004 *Phys. Rev Lett.* **92** 235004
[34] Tobita K. *et al* 2000 *JAERI-Conf. 2000-004 JAERI (Naka, Japan)*
[35] Kramer G.J. *et al* 2000 *Nucl. Fusion* **40** 1383
[36] Tobita K. *et al* 2002 *Fusion Sci. Technol.* **42** 315
[37] Graves J.P. 2004 *Phys. Rev Lett.* **92** 185003
[38] Günter S. *et al* 2003 *Nucl. Fusion* **43** 161
[39] Westerhof E. *et al* 2002 *Nucl. Fusion* **42** 1324
[40] Eriksson L.-G. *et al* 1998 *Phys. Rev Lett.* **81** 1231
[41] Mayoral M-L. *et al* 2002 *Proc. 29th EPS Conf. on Plasma Physics. and Controlled. Fusion (Montreux, Switzerland, 2002)* vol 26B (ECA) P-1.026
[42] Kolesnichenko Ya.I. *et al* 1992 *Phys. Rev. Lett.* **68** 3881
[43] Porcelli F. *et al* 2004 *Nucl. Fusion* **44** 362
[44] Bussac M.N. *et al* 1975 *Phys. Rev. Lett.* **35** 1638
[45] Kruskal M.D. and Oberman C.R. 1958 *Phys. Fluids* **1** 275
[46] Ward D.J. and Jardin S.C. 1989 *Nucl. Fusion* **29** 905
[47] Jardin S.C., Bell M.G. and Pomphrey N. 1993 *Nucl. Fusion* **33** 371
[48] Bateman G. *et al* 1998 *Phys. Plasmas* **5** 2355
[49] Waltz R. *et al* 1997 *Phys. Plasmas* **4** 2482
[50] Güenter S. *et al* 1998 *Nucl. Fusion* **38** 1431
[51] Maraschek M. *et al* 2003 *Plasma Phys. Control. Fusion* **45** 1369
[52] La Haye R.J. and Sauter O. 1998 *Nucl. Fusion* **38** 987
[53] Buttery R.J. *et al* 2003 *Nucl Fusion* **43** 69
[54] Sauter O. *et al* 2002 *Plasma Phys. Control. Fusion* **44** 1999
[55] Isayama A. *et al* 2001 *Nucl. Fusion* **41** 761
[56] Kislov D.A. *et al* 2001 *Nucl. Fusion* **41** 1619
[57] La Haye R.J. *et al* 2000 *Phys. Plasmas* **7** 3349
[58] Buttery R.J. *et al Proc. 20th Int. Conf. on Fusion Energy 2004 (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM EX/7-1 and http://www.naweb.iaea.org/napc/physics/fec/fec2004/index.html
[59] Glasser A.H. *et al* 1976 *Phys. Fluids* **19** 567
[60] Connor J.W *et al* 2001 *Phys. Plasmas* **8** 2835
[61] Fitzpatrick R. 1995 *Phys. Plasmas* **2** 825
[62] Mikhailovskii A.B. 2003 *Contrib. Plasma Phys.* **43** 125
[63] Reimerdes H. *et al* 2002 *Phys. Rev. Lett.* **88** 105005-1
[64] Brennan D.P. *et al* 2003 *Phys. Plasmas* **10** 1643
[65] Gude A., Guenter S. and Sesnic S. *et al* 1999 *Nucl. Fusion* **39** 127
[66] Frederickson E. *et al* 2002 *Phys. Plasma* **9** 548
[67] Gorelenkov N.N. *et al* 1996 *Phys. Plasmas* **3** 3379
[68] Konovalov S.V. *et al* 2002 *Plasma Phys. Control. Fusion* **44** L51
[69] Poli E. *et al* 2003 *Plasma Phys. Control. Fusion* **45** 71
[70] Buttery R.J. *et al* 2002 *Phys. Rev. Lett.* **88** 125005
[71] Lütjens H. *et al* 2002 *Phys. Plasmas* **9** 4837
[72] Waelbroeck F.L. *et al* 2001 *Phys. Rev. Lett.* **87** 215003-1
[73] Smolyakov A.I. *et al* 2004 *Plasma Phys. Control. Fusion* **46** L1-L6
[74] Smolyakov A.I. and Lazzaro E. 2004 *Phys. Plasmas* **11** 4353
[75] La Haye R.J. *et al* 2003 *Phys. Plasmas* **10** 3644
[76] Ozeki T. *et al* 2003 *Plasma Phys. Control. Fusion* **45** 645
[77] Bergmann A. *et al* 2002 *Proc. 19th Int. Conf. on Fusion Energy 2002 (Lyon, 2002)* (Vienna: IAEA) CD-ROM TH/P1-01 and http://www.iaea.org/programmes/ripc/physics/fec2002/html/fec2002.html
[78] Nave M.F.F. *et al* 2003 *Nucl. Fusion* **43** 179
[79] Itoh S.-I. *et al* 2004 *Plasma Phys. Control. Fusion* **46** 123
[80] Pustovitov V.D. 2005 *Nucl. Fusion* **45** 245
[81] Buttery R.J. *et al* 2001 *Proc. 28th EPS Conf. on Plasma Physics* (*Funchal, Portugal, 2001*) vol 25A (ECA) pp 1813–6 paper P-5.011 and http://www.cfn.ist.utl.pt/EPS2001/CD/pdfs/P5.011.pdf
[82] Popov A.M. *et al* 2002 *Phys. Plasmas* **9** 4205
[83] Hender T.C. *et al* 2004 *Nucl Fusion* **44** 788
[84] Hender T.C. *et al Proc. 18th Int. Conf. on Fusion Energy 2000 (Sorrento, 2000)* (Vienna: IAEA) CD-ROM EX/P3-02 and http://www.iaea.org/programmes/ripc/physics/fec2000/html/model.htm
[85] Kruger S.E. *et al* 1998 *Phys. Plasmas* **5** 455
[86] Mikhailovskii A.B. *et al* 2000 *Plasma Phys. Rep.* **26** 375
[87] Marchenko V.S. and Lutsenko V.V. 2001 *Phys. Plasmas* **8** 4834
[88] Zohm H. *et al* 2001 *Phys. Plasmas* **8** 2009
[89] Morris A.W. *et al* 1992 *Proc. 19th EPS Conf. on Controlled Fusion and Plasma Physics (Innsbruck, Austria, 1991)* vol 16C-I (ECA) p 423
[90] Warrick C.D. *et al* 2000 *Phys. Rev. Lett.* **85** 574
[91] Zohm H. *et al* 1999 *Nucl. Fusion* **39** 577
[92] Isayama A. *et al* 2000 *Plasma Phys. Control. Fusion* **42** L37
[93] La Haye R.J. *et al* 2002 *Phys. Plasmas* **9** 2051
[94] Petty C.C. *et al* 2004 *Nucl. Fusion* **44** 243

[95] Doyle E.J. *et al* 2007 Progress in the ITER Physics Basis *Nucl. Fusion* **42** S18–S127
[96] Zohm H. *et al* 2001 *Nucl. Fusion* **41** 197
[97] Maraschek M. *et al* 2004 *Nucl. Fusion* **45** 1369
[98] Gantenbein G. *et al* 2000 *Phys. Rev. Lett.* **85** 1242
[99] Isayama A. *et al* 2003 *Nucl. Fusion* **43** 1272
[100] Lazzaro E. *et al* 2005 *6th Int. Workshop on Strong Microwaves in Plasmas (Nizhny Novgorod)*
[101] Berrino J. *et al* 2006 *IEEE Trans. Nucl. Sci.* **53** 1009
[102] Nagasaki K. *et al* 2003 *Nucl. Fusion* **43** L7
[103] Isayama A. *et al* 2005 *Phys. Plasmas* **12** 056117
[104] Buttery R.J. *et al* 2004 *Nucl. Fusion* **44** 678
[105] Wolf R.C. *et al* 1999 *Plasma Phys. Control. Fusion* **41** B93
[106] La Haye R.J., Rice B.W. and Strait E.J. 2000 *Nucl. Fusion* **40** 53
[107] Campbell D.J. *et al* 1988 *Phys. Rev. Lett.* **60** 2148
[108] Maget P. *et al* 2002 *Theory of Fusion Plasmas Proc. Joint Varenna–Lausanne Int. Workshop (Varenna 2002)* ed J.W. Connor *et al* (Bologna: Editrice Compositori) p 363 ISPP-20
[109] Stober J. *et al* 2002 *Plasma Phys. Control. Fusion* **44** A159
[110] Angioni C. *et al* 2003 *Phys. Plasmas* **10** 3225
[111] Garbet X. *et al* 2003 *Phys. Rev. Lett.* **91** 035001
[112] Angioni C. *et al* 2004 *Nucl. Fusion* **44** 827
[113] Gude A. *et al* 2002 *Nucl. Fusion* **42** 833
[114] Günter S. *et al* 2001 *Phys. Rev. Lett.* **87** 275001
[115] Günter S. *et al* 2004 *Nucl. Fusion* **44** 524
[116] Sips A.C.C. *et al* 2002 *Plasma Phys. Control. Fusion* **44** B69
[117] Raju D., Sauter O. and Lister J.B. 2003 *Plasma Phys. Control. Fusion* **45** 369
[118] Yu Q., Günter S. and Lackner K. 2000 *Phys. Rev. Lett* **85** 2949
[119] St. John H.E. *et al* 1995 *Proc. 15th Int. Conf. on Plasma Physics and Controlled Nuclear Fusion Research 1994 (Seville, 1994)* vol 3 p 603 (Vienna: IAEA)
[120] Lao L.L. *et al* 1990 *Nucl. Fusion* **30** 1035
[121] Matsuda K. 1989 *IEEE Trans. Plasma Sci.* **PS-17** 6
[122] Prater R. *et al* 2003 *Nucl. Fusion* **43** 1128
[123] La Haye R.J. *et al* 2006 *Nucl. Fusion* **46** 451
[124] Harvey R.W. and Perkins F.W. 2001 *Nucl. Fusion* **41** 1847
[125] Hamamatsu K. and Fukuyama A. 2000 *Plasma Phys. Control. Fusion* **42** 1309
[126] Pustovitov V.D. *et al* 2000 *Proc. 18th Int. Conf. on Fusion Energy 2000 (Sorrento, Italy, 2000)* (Vienna: IAEA) CD-ROM ITERP/07 and http://www.iaea.org/programmes/ripc/physics/fec2000/html/nodel.htm
[127] Hayashi N. *et al* 2004 *Nucl. Fusion* **44** 477
[128] Hayashi N. *et al* 2004 *J. Plasma Fusion Res.* **80** 605
[129] Hegna C.C. and Callen J.D. 1997 *Phys. Plasmas* **4** 2940
[130] Strait E.J. *et al* 1995 *Phys. Rev. Lett.* **74** 2483
[131] Taylor T.S. *et al* 1995 *Phys. Plasmas* **2** 2390
[132] Bondeson A. and Ward D.J. 1994 *Phys. Rev. Lett.* **72** 2709
[133] Chu M.S. *et al* 1995 *Phys. Plasmas* **2** 2236
[134] Bondeson A. and Chu M.S. 1996 *Phys. Plasmas* **3** 3013
[135] Kuvshinov B.N. Mikhailovskii A.B. 1998 *Plasma Phys. Rep.* **24** 623
[136] Garofalo A.M. *et al* 2000 *Proc. 18th Int. Conf. on Fusion Energy 2000 (Sorrento, Italy, 2000)* (Vienna: IAEA) CD-ROM EXP3-01 and http://www.iaea.org/programmes/ripc/physics/fec2000/html/nodel.htm
[137] Boozer A.H. 2001 *Phys. Rev. Lett.* **86** 5059
[138] Garofalo A.M., Jensen T.H. and Strait E.J. 2002 *Phys. Plasmas* **9** 4573
[139] Garofalo A.M. *et al* 2002 *Phys. Rev. Lett.* **89** 235001
[140] La Haye R.J. *et al* 2005 *Nucl. Fusion* **44** 1197
[141] Hender T.C. *et al* 2004 *Proc. 20th Int. Conf. on Fusion Energy (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM EX/P2-22 and http://www.naweb.iaea.org/napc/physics/fec/fec2004/index.html
[142] Reimerdes H. *et al* 2006 *Phys. Plasmas* **13** 056107
[143] Pustovitov V.D. 2003 *JETP Lett.* **78** 281
[144] Gregoratto D. *et al* 2001 *Plasma Phys. Control. Fusion* **43** 1425
[145] Bondeson A. *et al* 2003 *Plasma Phys. Control. Fusion* **45** A253
[146] Liu Y.Q. *et al* 2004 *Nucl. Fusion* **44** 232
[147] Kurita G. *et al* 2003 *Nucl. Fusion* **43** 949
[148] Boozer A.H. and Maslovsky D. 2004 *Proc. 31st EPS Conf. on Plasma Physics* (*London, UK, 2004*) vol 28B (ECA) P-2.159
[149] Okabayashi M. *et al* 2002 *J. Plasma Fusion Res. Ser.* **5** 42
[150] Okabayashi M. *et al* 2002 *Plasma Phys. Control. Fusion* **44** B339
[151] Gryaznevich M.P. *et al* 2003 *APS (Albuquerque, New Mexico, 2003)* paper RP1.036 and http://www.aps.org/meet/DPP03/baps/abs/S2080036.html
[152] Liu YQ. *et al Proc. 20th Int. Conf. on Fusion Energy 2004 (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM TH2-1 and http://www.naweb.iaea.org/napc/physics/fec/fec2004/index.html
[153] Garofalo A.M., Jensen T.H. and Strait E.J. 2003 *Phys. Plasmas* **10** 4776
[154] Garofalo A.M. *et al* 2006 *Proc. 21st Int. Conf. on Fusion Energy Research 2006 (Chengdu, China, 2006)* (Vienna: IAEA) CD-ROM EX/7-1Ra
[155] Takechi M. *et al* 2006 *Proc. 21st Int. Conf. on Fusion Energy Research 2006 (Chengdu, China, 2006)* (Vienna: IAEA) CD-ROM EX/7-1Rb
[156] Polevoi A.R. *et al* 2002 *J. Plasma Fusion Res. Ser.* **5** 82
[157] Liu Y.Q. *et al* 2000 *Phys. Plasmas* **7** 3681
[158] Bondeson A. *et al* 2002 *Nucl. Fusion* **42** 768
[159] Okabayashi M., Pomphrey N. and Hatcher R.E. 1998 *Nucl. Fusion* **38** 1607
[160] Pustovitov V.D. 2001 *Plasma Phys. Rep.* **27** 195
[161] Strait E.J. *et al* 2003 *Nucl. Fusion* **43** 430
[162] Okabayashi M. *et al* 2001 *Phys. Plasmas* **8** 2071 075001
[163] Pustovitov V.D. 2002 *J. Plasma Fusion Res. Ser.* **5** 278
[164] Bialek J. *et al* 2001 *Phys. Plasmas* **8** 2170
[165] Chu M.S. *et al* 2003 *Nucl. Fusion* **43** 441
[166] Fitzpatrick R. 2001 *Phys. Plasmas* **8** 871
[167] Gribov Y. and Pustovitov V.D. 2002 *Proc. 19th Int. Conf. on Fusion Energy 2002 (Lyon, France, 2002)* (Vienna: IAEA), CD-ROM CT/P-12 and http://www.iaea.org/programmes/ripc/physics/fec2002/html/fec2002.htm
[168] Liu Y.Q. and Bondeson A. 2002 *Plasma Phys. Control. Fusion* **44** L21
[169] Chu M.S. *et al* 2004 *Phys. Plasmas* **11** 2497
[170] Liu Y.Q. *et al* 2004 *Nucl. Fusion* **44** 77
[171] Medvedev S.Yu. and Pustovitov V.D. 2003 *Plasma Phys. Rep.* **29** 1009
[172] Okabayashi M. *et al* 1996 *Nucl. Fusion* **36** 1167
[173] Sabbagh S.A. *et al* 2002 *Phys. Plasmas* **9** 2085
[174] Takeji S. *et al* 2002 *Nucl. Fusion* **42** 5
[175] Pinches S.D. *et al* 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St Petersburg, Russia, 2003)* vol 27A (ECA) P-1.93
[176] Reimerdes H. *et al* 2005 *Proc. 32nd EPS Conf. on Plasma Physics and Controlled Fusion (Tarragona, Spain, 2005)* paper P5.056
[177] Garofalo A.M. *et al* 2001 *Nucl. Fusion* **41** 1171
[178] Boozer A.H. 2003 *Phys. Plasmas* **10** 1458
[179] Pustovitov V.D. 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St. Petersburg, Russia, 2003)* vol 27A (ECA) P-4.167 and http://eps2003.ioffe.ru/
[180] Pustovitov V.D. 2002 *Plasma Phys. Control. Fusion* **44** 295
[181] Garofalo A.M., *et al* 1999 *Phys. Rev. Lett.* **82** 3811
[182] Takeji S. *et al* 2002 *J. Plasma Fusion Res.* **78** 447
[183] Garofalo A.M. *et al* 2002 *Phys. Plasmas* **9** 1997
[184] Shilov M. *et al* 2004 *Phys. Plasmas* **11** 2573
[185] Sabbagh S.A. *et al* 2000 *Proc. 20th Int. Conf. on Fusion Energy 2004 (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM EX/3-2 and http://www.naweb.iaea.org/napc/physics/fec/fec2004/index.html
[186] Garofalo A.M., La Haye R.J. and Scoville J.T. 2002 *Nucl. Fusion* **42** 1335

[187] Wade M.R. *et al* 2003 *Nucl. Fusion* **43** 634
[188] Anderson P.M. *et al* 2003 *Fusion Eng. Des.* **66–68** 791
[189] Jackson G.L. *et al* 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St Petersburg, Russia, 2003)* vol 27A (ECA) P-4.47
[190] Strait E.J. *et al* 2004 *Phys. Plasmas* **11** 2505
[191] Menard J.E. *et al* 2003 *Nucl. Fusion* **43** 330
[192] Sabbagh S.A. *et al* 2004 *Nucl. Fusion* **44** 560
[193] Tsuzuki K. *et al* 2004 *Proc. 20th Int. Conf. on Fusion Energy 2004 (Vilamoura, Portugal, 2004)*
[194] Kurita G. *et al* 2004 *Proc. 20th Int. Conf. on Fusion Energy 2004 (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM FT/P7-7 and http://www.naweb.iaea.org/napc/physics/fec/fec2004/index.html
[195] Ward D.J. and Bondeson A. 1995 *Phys. Plasmas* **2** 1570
[196] Bondeson A. *et al* 2002 *Phys. Plasmas* **9** 2044
[197] Cates C. *et al* 2000 *Phys. Plasmas* **7** 3133
[198] Maurer D.A. *et al Proc. 19th Int. Conf. on Fusion Energy 2002 (Lyon, France, 2002)* (Vienna: IAEA) CD-ROM TH/P3-13 and http://www.iaea.org/programmes/ripc/physics/fec2002/html/fec2002.htm
[199] Bondeson A. *et al* 2001 *Nucl. Fusion* **41** 455
[200] Boozer A.H. 1998 *Phys. Plasmas* **5** 3350
[201] Reimerdes H. *et al* 2004 *Phys. Rev. Lett.* **93** 135002
[202] Reimerdes H. *et al* 2004 *Proc. 20th Int. Conf. on Fusion Energy 2004 (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM EX/3-1Rb and http://www.naweb.iaea.org/napc/physics/fec//fec2004/index.html
[203] Polevoi A.R. *et al Proc. 19th Int. Conf. on Fusion Energy 2002 (Lyon, France, 2002)* (Vienna: IAEA) CD-ROM CT/P-08 and http://www.iaea.org/programmes/ripc/physics/fec2002/html/fec2002.htm
[204] Navratil G.A. *et al* 2003 RWM Control in FIRE and ITER *Workshop on Active Control of MHD stability (University of Texas-Austin)*
[205] Pustovitov V.D. and Medvedev S. Yu. 2003 Modeling of feedback stabilization of RWM in the T-15M tokamak *Web Proc. EPS 2003 (St Petersburg, July 2003)* paper P-3.136 and http://eps2003.ioffe.ru/
[206] Liu Y.Q. *et al* 2005 *Nucl. Fusion* **45** 1131
[207] Medvedev S.Yu. 2003 private communication, (Moscow State University)
[208] Bialek J. 2005 private communication (Columbia University, New York)
[209] Manickam J. 2003 private communication (PPPL, USA)
[210] Buttery R.J. *et al* 2000 *Nucl. Fusion* **40** 807
[211] La Haye R.J. *et al* 1992 *Phys. Fluids* B **4** 2098
[212] Buttery R.J. *et al* 1999 *Nucl. Fusion* **39** 1827
[213] Scoville J.T. and La Haye R.J. 1997 *Bull. Am. Phys. Soc.* **42** 1979
[214] Buttery R.J. *et al* 1997 *Proc. 24th EPS Conf. on Controlled Fusion and Plasma Physics (Berchtesgaden, Germany, 1997)* vol 21A 265
[215] Scoville J.T. and La Haye R.J. 1996 *Bull. Am. Phys. Soc.* **41** 1570
[216] Wolfe S. *et al* 2005 *Phys. Plasmas* **12** 056110
[217] Koslowski H.R. *et al* 2004 *Proc. 31st EPS Conf. on Plasma Physics (London, UK, 2004)* paper P1.124
[218] Lazzaro E. *et al* 2002 *Phys. Plasmas* **9** 3906
[219] Fitzpatrick R. 1993 *Nucl. Fusion* **33** 1049
[220] Lazzaro E. and Zanca P. 2003 *Phys. Plasmas* **10** 2399
[221] Hender T.C. *et al* 2002 *Proc. 19th Int. Conf. on Fusion Energy 2002 (Lyon, France, 2002)* (Vienna: IAEA) CD-ROM EX/S1-2 and http://www.iaea.org/programmes/ripc/physics/fec2002/html/fec2002.htm
[222] Shaing K.C. 2003 *Phys. Plasmas* **10** 1443
[223] La Haye R.J *et al* 1992 *Nucl. Fusion* **32** 2119
[224] Zhu W. *et al* 2006, *Phys Rev Lett* **96** 225002
[225] Fitzpatrick R. and Hender T.C. 1991 *Phys. Fluids* B **3** 644
[226] Scoville J.T. *et al* 2003 *45th Annual Meeting of the Division of Plasma Physics, American Physical Society (Albuquerque, USA, 2003) Bull. Am. Phys. Soc.* **48** 264 paper QP1.074 and http://www.aps.org/meet/DPP03/baps/abs/S2080036.html
[227] Scoville J.T. *et al* 2004 *46th Annual Meeting of the Division of Plasma Physics, American Physical Society (Savannah, USA, 2004)* paper NP1.008
[228] Luxon J.L. *et al* 2003 *Nucl. Fusion* **43** 1813
[229] Schaffer M.J *et al* 2003 *45th Annual Meeting of the Division of Plasma Physics, American Physical Society (Albuquerque, USA, 2003) Bull. Am. Phys. Soc.* **48** 264, poster QP1.035 and http://www.aps.org/meet/DPP03/baps/abs/S2080036.html
[230] Scoville J.T. and La Haye R.J. 2003 *Nucl. Fusion* **43** 250
[231] Isei N. *et al* 2001 *Fusion Technol.* **39** 1101
[232] Brennan D.P. *et al* 2002 *Phys. Plasmas* **9** 2998
[233] Amoskov V. *et al* 2004 *Plasma Devices Oper.* **12** 285
[234] Amoskov V. *et al* 2005 *Plasma Devices Oper.* **13** 87
[235] Lin-Liu Y.R. *et al* 1999 *Phys. Plasmas* **6** 3934
[236] Gruber O. *et al* 1999 *Phys. Rev. Lett.* **83** 1787
[237] Kamada Y. *et al* 1999 *Nucl. Fusion* **39** 1845
[238] Rice B.W. *et al* 1999 *Nucl. Fusion* **39** 1855
[239] Wade M.R. *et al* 2001 *Phys. Plasmas* **8** 2208
[240] Joffrin E. *et al* 2005 *Nucl. Fusion* **45** 626
[241] Wade M.R. *et al* 2005 *Nucl. Fusion* **45** 407
[242] Staebler A. *et al* 2005 *Nucl. Fusion* **45** 617
[243] Lao L.L. *et al* 1996 *Phys. Plasmas* **3** 1951
[244] Takeji S. *et al* 1997 *Phys. Plasmas* **4** 4283
[245] Huysmans G.T.A. *et al* 1999 *Nucl. Fusion* **39** 1489
[246] Günter S. *et al* 2001 *Nucl. Fusion* **41** 1283
[247] Manickam J. *et al* 1999 *Nucl. Fusion* **39** 1819
[248] Turnbull A.D. *et al* 1998 *Nucl. Fusion* **38** 1467
[249] Sarazin Y. *et al* 2002 *Plasma Phys. Control. Fusion* **44** 2445
[250] Burrell K.H. *et al* 2001 *Phys. Plasmas* **8** 2153
[251] Doyle E.J. *et al* 2002 *Nucl. Fusion* **42** 333
[252] Greenfield C.M. *et al* 2002 *Plasma Phys. Control. Fusion* **44** A123
[253] Menard J.E. *et al* 2004 *Phys. Plasmas* **11** 639
[254] Luce T.C. *et al* 2001 *Nucl. Fusion* **41** 1585
[255] Ferron J.R. *et al* 2005 *Phys. Plasmas* **12** 056126
[256] Buttery R.J. *et al* 2004 *Nucl. Fusion* **44** 1027
[257] Takeji S. *et al* 2000 *J. Plasma Fusion Res.* **76** 575
[258] Hender T.C. *et al* 2002 *Plasma Phys. Control. Fusion* **44** 1143
[259] Turnbull A.D. *et al* 2002 *Nucl. Fusion* **42** 917
[260] Günter S. *et al* 2000 *Nucl. Fusion* **40** 1541
[261] Callen J.D. *et al* 1999 *Phys. Plasmas* **6** 2963
[262] Turnbull A.D. *et al* 1997 *Proc. 16th Int. Conf. on Fusion Energy 1996 (Montreal, Canada, 1996)* vol 2 (Vienna: IAEA) p 509
[263] Ishii Y. *et al* 1998 *Plasma Phys. Control. Fusion* **40** 1607
[264] Ozeki T. *et al* 1995 *Nucl. Fusion* **35** 861
[265] Makowski M.A. *et al* 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St Petersburg, Russia, 2003)* vol P-2 p 113 and http://eps2003.ioffe.ru/
[266] Greenfield C.M. *et al* 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St Petersburg, Russia, 2003)* vol P-4 p 92 and http://eps2003.ioffe.ru/
[267] Kessel C.E. *et al* 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St Petersburg, Russia, 2003)* vol P-4 p 44 and http://eps2003.ioffe.ru/
[268] Ishii Y. *et al* 2000 *Phys. Plasmas* **7** 4477
[269] Popov A.M. *et al* 2001 *Phys. Plasmas* **8** 3605
[270] Fujita T. *et al* 2001 *Phys. Rev. Lett.* **87** 245001
[271] Hawkes N.C. *et al* 2001 *Phys. Rev. Lett.* **87** 115001
[272] Chu M.S. and Parks P. 2002 *Phys. Plasmas* **9** 5036
[273] Takizuka T. 2002 *J. Plasma Fusion Res.* **78** 1282
[274] Martynov A.A. *et al* 2003 *Phys. Rev. Lett.* **91** 085004
[275] Huysmans G.T.A. *et al* 2001 *Phys. Rev. Lett.* **87** 245002
[276] Huysmans G.T.A. *et al* 2002 *Proc. 29th EPS Conf. on Plasma Physics and Controlled Fusion (Montreux, Switzerland, 2002)* pp 17–21

[277] Stratton B.C. *et al* 2002 *Plasma Phys. Control. Fusion* **44** 1127
[278] Ozeki T. *et al* 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St Petersburg, Russia, 2003)* vol P-2 pp 111 and http://eps2003.ioffe.ru/
[279] Takeji S. *et al* 2002 *Fusion Sci. Technol.* **42** 278
[280] Shimomura Y. *et al* 2001 *Plasma Phys. Control. Fusion* **43** A385
[281] Holties H.A. *et al* 1996 *Nucl. Fusion* **36** 973
[282] Bondeson A. *et al* 1997 *Nucl. Fusion* **37** 1419
[283] Pereverzev G.V. *et al* 1991 *Report IPP*5/42
[284] Degtyarev L. *et al* 1997 *Comput. Phys. Commun.* **103** 10
[285] Polevoi A.R. *et al Proc. 20th Int. Conf. on Fusion Energy 2004 (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM IT/P3-28 and http://www.naweb.iaea.org/napc/physics/fec/fec2004/index.html
[286] Shimada M. *et al* 2005 *Nucl. Fusion* **44** 350
[287] Parker R.R 2000 *Nucl. Fusion* **40** 473
[288] Aymar R. *et al* 2002 *Proc. 19th Int. Conf. on Fusion Energy 2002 (Lyon, France, 2002)* (Vienna: IAEA) CD-ROM OV/1-1 and http://www.iaea.org/programmes/ripc/physics/fec2002/html/fec2002.htm
[289] Sokolov Yu. A. 1979 *JETP Lett.* **29** 244–6
[290] Rosenbluth M.N. and Putvinski S.V. 1997 *Nucl. Fusion* **37** 1355
[291] Plyusnin V.V. *et al* 2006 *Nucl. Fusion* **46** 277 *Plasma Physics (Funchal, Portugal, 2001)* vol 25A (ECA) pp 1805
[292] Kadomtsev B.B. 1975 *Plasma Phys. (Russ)* **1** 710
[293] Kadomtsev B.B. and Pogutse O. 1974 *Sov. Phys.–JETP* **38** 283
[294] Wesson J.A. *et al* 1997 *Phys. Rev. Lett.* **79** 5018
[295] Taylor P.L. *et al* 1996 *Phys. Rev. Lett.* **76** 916
[296] Helander P. *et al* 2002 *Phys. Rev. Lett.* **89** 235002
[297] Mirnov S. *et al* 1998 *Phys. Plasmas* **5** 3950
[298] Mirnov S.V. 2001 *Proc. 28th EPS Conf. on Controlled Fusion and Plasma Physics (Funchal, Portugal, 2001)* vol 25A (ECA) pp 1473–6
[299] Kleva R.G. and Guzdar P.N. 2001 *Phys. Plasmas* **8** 103
[300] Kruger S.E. *et al* 2004 *Computer Phys. Commun.* **164** 34
[301] Cowley S.C. *et al* 2003 *Plasma Phys. Control. Fusion* **45** A31
[302] Salzedas F. *et al* 2003 *Phys. Plasmas* **9** 3402
[303] Salzedas F. *et al* 2002 *Proc. 29th EPS Conf. on Controlled Fusion and Plasma Physics (Montreux, Switzerland, 2002)* P-1.039
[304] Salzedas F. *et al* 2003 *Proc. 30th EPS Conf. Controlled Fusion and Plasma Physics (St Petersburg, Russia 2003)* vol 27A P-2.95 and http://eps2003.ioffe.ru/
[305] Hyatt A.W. *et al* 2001 *Bull. Am. Phys Soc.* **45** 300
[306] Pautasso G. *et al* 2001 *J. Nucl. Mater.* **290–293** 1045
[307] 2001 *ITER Final Design Report* complied by the ITER Director (http://www.naka.jaea.go.jp/ITER/FDR/)
[308] Riccardo V. *et al* 2005 *Nucl. Fusion* **45** 1427
[309] Schuller F.C. *et al* 1995 *Plasma Phys. Control. Fusion* **37** A135
[310] Pautasso G. *et al* 2003 *Proc. 30th EPS Conf. on Plasma Physics (St Petersburg, Russia, 2003)* vol 27A P-1.135 and http://eps2003.ioffe.ru/
[311] Pautasso G. *et al* 2004 *Proc. 31st EPS Conf. on Plasma Physics (London, UK, 2004)* **28G** P-4.132
[312] Konz C. *et al* 2005 *Proc. 32nd EPS Conf. on Plasma Physics and Controlled Fusion, (Tarragona, Spain, 2005)* vol 29C pp O2-005
[313] Counsell G.F. *et al* 2005 *Nucl. Fusion* **45** S157
[314] Andrew P. *et al* 2006 Main chamber power loads during disruptions *Preprint* EFD-C(06)02/08
[315] Herrmann A. *et al* 2004 *Plasma Phys. Control. Fusion* **46** 971
[316] Ciotti M. *et al* 1999 *J. Nucl. Mater.* **266–269** 1023
[317] Finken K.H. *et al* 2001 *J. Nucl. Mater.* **290–293** 1064
[318] Wesley J. *et al* 1998 *IEEE Proc. 17th Symp. Fusion Engineering (San Diego, USA, 1997)* vol 1 (Piscataway: IEEE) p 483
[319] Tamai H. *et al* 2002 *Nucl. Fusion* **42** 290
[320] Pautasso G. and Gruber O. *Fusion Sci. Technol.* **44** (2003) 716
[321] Neyatani Y. *et al Proc. 17th Int. Conf. on Fusion Energy 1998 (Yokohama, Japan, 1998)* (Vienna: IAEA) CD-ROM file EXP3/11 and http://www.iaea.org/programmes/ripc/physics/start.htm
[322] Andrew P. *et al* 2005 *J. Nucl. Mater.* **337–339** 99
[323] Andrew P. *et al* 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St Petersburg, Russia, 2003)* vol 27A P-1.108 and http://eps2003.ioffe.ru/
[324] Whyte D.G. *et al* 2003 *J. Nucl. Mater.* **313–316** 1239
[325] Loarte A. *et al* 2004 *Proc. 20th Int. Conf. on Fusion Energy 2004 (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM IT/P3-34 and http://www.naweb.iaea.org/napc/physics/fec/fec2004/index.html
[326] Sugihara M. *et al* 2003 *Plasma Fusion Res.* **79** 706
[327] Sugihara M. *et al* 2004 *Proc. 20th Int. Conf. on Fusion Energy 2004 (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM IT/P3-29 and http://www-naweb.iaea.org/napc/physics/fec/fec2004/datasets/index.html
[328] Hyatt A. 2004 private communication (San Diego: General Atomics)
[329] Wesley J. *et al* 2006 *Proc. 21st Int. Conf. on Fusion Energy Research 2006 (Chengdu, China, 2006)* (Vienna: IAEA) CD-ROM IT/P1-21
[330] Nakamura Y. *et al* 2002 *Plasma Phys. Control. Fusion* **44** 1471
[331] Humphreys D.A. and Kellman A.G. 1999 *Phys. Plasmas* **6** 2742
[332] Neyatani Y. *et al* 1999 *Nucl. Fusion* **39** 559
[333] Riccardo V. *et al* 2004 *Plasma Phys. Control. Fusion* **46** 925
[334] Yoshino R., Nakamura Y. and Neyatani Y. 1996 *Nucl. Fusion* **36** 295
[335] Riccardo V. *et al* 2003 *Plasma Phys. Control. Fusion* **45** A269
[336] Counsell G.F. *et al* 2003 *Nucl. Fusion* **43** 1197
[337] Irby J. *et al* 2006 *Fusion Sci. Technol.* **51** 460
[338] Khayrutdinov R.R. and Lukash V.E. 1993 *Comput. Phys. Commun.* **109** 193
[339] Riccardo V., Walker S. and Noll P. 2000 *Plasma Phys. Control. Fusion* **42** 29
[340] Riccardo V. Walker S. and Noll P. 2000 *Fusion Eng. Des.* **47** 389
[341] Riccardo V., Noll P. and Walker S.P. 2000 *Nucl. Fusion* **40** 1805
[342] Harvey R.W. *et al* 2000 *Phys. Plasmas* **7** 4590
[343] Yoshino R., Tokuda S. and Kawano Y. 1999 *Nucl. Fusion* **39** 151
[344] Gill R.D. *et al* 2002 *Nucl. Fusion* **42** 1039
[345] Martin G. 2000 *Proc. 6th IAEA TCM on Energetic Particles in Magnetic Confinement Systems (Naka, Japan, 2000)* JAERI-Conf 2000-004, 1
[346] Gill R.D. *et al* 2000 *Nucl. Fusion* **40** 163
[347] Helander P., Eriksson L.-G. and Andersson F. 2002 *Plasma Phys. Control. Fusion* **44** B247
[348] Helander P. *et al* 2003 Runaway electrons and current dynamics during tokamak disruptions *8th IAEA Technical Meeting on Energetic Particles in Magnetic Confinement Systems (San Diego, 6–8 October 2003)*
[349] Eriksson L.-G. and Helander P. 2003 *Comput. Phys. Commun.* **154** 175
[350] Schittenhelm M. 1997 *Proc. 24th Conf. on Controlled Fusion and Plasma Physics (Berchtesgarden, Germany, 1997)* vol 21A, Part III, (Geneva: European Physical Society) p 985
[351] Martin G. *et al* 2004 *Proc. 20th Int. Conf. on Fusion Energy (Vilamoura, Portugal, 2004)* (Vienna: IAEA) CD-ROM EX/10-6Rc and http://www-naweb.iaea.org/napc/physics/fec/fec2004/datasets/index.html
[352] Lukash V.E. and Khayrutdinov R.R. 2000 *Proc. 6th IAEA TCM on Energetic Particles in Magnetic Confinement Systems (Naka, Japan, 2000)* JAERI-Conf 2000-004, 13
[353] Maddaluno G. *et al* 2003 *J. Nucl. Mater.* **313–316** 651

[354] Yoshino R., Nakamura Y. and Neyatani Y. 1997 *Nucl. Fusion* **37** 1161
[355] Andersson F., Helander P. and Eriksson L.-G. 2001 *Phys. Plasmas* **8** 5221
[356] Kawano Y. *et al* 2005 *J. Plasma Fusion Res.* **81** 593
[357] Yoshino R. and Tokuda S. 2000 *Nucl. Fusion* **40** 1293
[358] Kawano Y. *et al* 2005 *J. Plasma Fusion Res.* **81** 743
[359] Tokuda S. and Yoshino R. 1999 *Nucl. Fusion* **39** 1123
[360] Martin-Solis J.R. *et al* 2006 *33rd EPS Conf. Plasma Physics (Rome, Italy, 2006)* P5.078
[361] Kawano Y. *et al* 1997 *Proc. 16th Int. Conf. on Fusion Energy 1996 (Montreal, Canada, 1996)* vol 1 (Vienna: IAEA) p 345
[362] Kawano Y. *et al* 2002 *Fusion Sci. Technol.* **42** 298
[363] Helander P., Eriksson L.-G. and Andersson F. 2000 *Phys. Plasmas* **7** 4106
[364] Jardin S.C. *et al* 2000 *Nucl. Fusion* **40** 923
[365] Taylor P.L. *et al* 1999 *Phys. Plasmas* **6** 1872
[366] Whyte D.G. *et al* 2002 *Phys. Rev. Lett.* **89** 055001
[367] Bakhtiari M. *et al* 2002 *Nucl. Fusion* **42** 1197
[368] Bakhtiari M. *et al* 2005 *Nucl. Fusion* **45** 318
[369] Humphreys D.A. and Whyte D.G. 2000 *Phys. Plasmas* **7** 4057
[370] Whyte D.G., Humphreys D.A. and Taylor P.L. 2000 *Phys. Plasmas* **7** 4052
[371] Hollmann E.M. *et al* 2005 *Nucl. Fusion* **45** 1046
[372] Kawano Y. *et al* 2001 Study on characteristics of runway electrons in JT-60U *18th JSPF Meeting (Fukuoka, 27–30 November 2001)*
[373] Gribov Y. *et al* 2007 Progress in the ITER Physics Basis *Nucl. Fusion* **47** S385–S403
[374] Lukash V.E. and Khayrutdinov R.R. 1996 *Plasma Phys. Rep.* **22** 91
[375] Jardin S.C. *et al* 1986 *Comput. Phys. Commun.* **66** 481
[376] Sugihara M. *et al* 2004 *Plasma Phys. Control. Fusion* **46** 1581
[377] Humphreys D.A. *et al* 1997 *General Atomics Report* GA-22692
[378] Khayrutdinov R.R. *et al* 2001 *Plasma Phys. Control. Fusion* **43** 321
[379] Lukash V.E. 2002 Validation of DINA halo area expansion model against JT-60U disruption data *2nd Meeting of the ITPA Topical Group on MHD, Disruptions and Control (Garching, Germany, October 2002)*
[380] Park W. *et al* 1999 *Phys. Plasmas* **6** 1796
[381] Paccagnella R. *et al* 2005 *Fusion Eng. Des.* **75–79** 589
[382] Pletzer A. http://w3.pppl.gov/rib/repositories/NTCC/catalog/Asset/grin.html
[383] Paccagnella R. *et al* 2003 Halo current simulation for tokamak plasmas *45th APS Division Plasma Physics Meeting (Albuquerque, USA, 2003)* paper QP1.079 and http://www.aps.org/meet/DPP03/baps/abs/S2080036.html
[384] Wroblewski D., Jahns G.L. and Leuer J.A. 1997 *Nucl. Fusion* **37** 725
[385] Sengupta A. and Ranjan P. 2001 *Nucl. Fusion* **41** 487
[386] Pautasso G. *et al* 2002 *Nucl. Fusion* **42** 100
[387] Yoshino R. 2003 *Nucl. Fusion* **43** 1771
[388] Yoshino R. 2005 *Nucl. Fusion* **45** 1232
[389] Cannas B. *et al* 2004 *Nucl. Fusion* **44** 68
[390] Windsor C.G. *et al* 2005 *Nucl. Fusion* **45** 337
[391] Fishpool G.M. and Haynes P.S. 1994 *Nucl. Fusion* **34** 109
[392] La Haye R.J. *et al* 1997 *General Atomics Report* GA-A22468
[393] Mertens V. *et al* 2003 *Fusion Sci. Technol.* **44** 593
[394] Esser H.G. *et al* 1997 *J. Nucl. Mater.* **241–243** 861
[395] Martin Y. *et al* 1998 *Proc. 25th EPS Conf. on Controlled Fusion and Plasma Physics (Prague, Czech Republic, 1998)* P3.017
[396] Scoville J. *et al* 1991 *Nucl. Fusion* **31** 875
[397] Kraemer-Flecken A. *et al* 2003 *Nucl. Fusion* **43** 1437
[398] Kraemer-Flecken A. *et al* 2001 *Fusion Eng. Des.* **56–57** 773
[399] Hoshino K. *et al* 1995 *Proc. 15th Int. Conf. on Plasma Physics and Controlled Nuclear Fusion Research 1994 (Seville, 1994)* (Vienna: IAEA) vol 1 p 697
[400] Salzedas F. *et al* 1999 *Proc. 26th EPS Conf. on Controlled Fusion and Plasma Physics (Maastricht, The Netherlands, 1999)* vol 23J p 625
[401] Yoshino R. *et al* 1994 *J. Plasma Fusion Res.* **70** 1081
[402] Oikawa T. *et al* 2004 *Fusion Eng. Des.* **70** 175
[403] Joffrin E. *et al* 2003 *Plasma Phys. Control. Fusion* **45** A367
[404] Bucalossi J. *et al* 2001 *J. Nucl. Mater.* **290–293** 566
[405] Reimerdes H. *et al* 2003 *Proc. 30th EPS Conf. on Controlled Fusion and Plasma Physics (St Petersburg, Russia, 2003)* vol 27A P-4.45 and http://eps2003.ioffe.ru/
[406] Testa D. *et al* 2000 *Proc. 27th EPS Conf. on Controlled Fusion and Plasma Physics (Budapest, Hungary, 2000)* vol 24B p 1429 and (http://202.127.204.25/gjhy/EPS/27th(2000)/pdf/p4_044.pdf) (2004) (Vienna: IAEA) CD-ROM EX/P2-33 and http://www-naweb.iaea.org/napc/physics/fec/fec2004/index.html
[407] Franzen P. *et al* 1998 *Fusion Technol.* **33** 84
[408] Mertens V. *et al* 2003 *Fusion Eng. Des.* **66–68** 119
[409] Granetz R.S. *et al* 1996 *Proc. 16th Int. Conf. on Fusion Energy 1996 (Montreal, Canada, 1996)* vol 1 (Vienna: IAEA) p 757
[410] Pautasso G. *et al* 1996 *Nucl. Fusion* **36** 1291
[411] Strauss H.R. and Park W. 1998 *Phys. Plasmas* **5** 2676
[412] Whyte D.G. *et al* 1998 *Phys. Rev. Lett.* **81** 4392
[413] Rosenbluth M.N., Putvinskij S.V. and Parks P.B. 1997 *Nucl. Fusion* **37** 955
[414] Pautasso G. *et al* 2002 *Proc. 29th EPS Conf. on Plasma Physics and Controlled Fusion (Montreux, Switzerland, 2002)* vol 26B, P2.051
[415] Finken K.H. *et al* 2001 *Nucl. Fusion* **41** 1651
[416] Jernigan T.C. *et al* 2005 private communication (ORNL USA)
[417] Izzo V.A. 2006 *Nucl. Fusion* **46** 541
[418] Parks P.B. *et al* 1997 *Fusion Technol.* **35** 267
[419] Summers D.A. 1995 *Waterjetting Technology* (London: Chapman & Hall)
[420] Nagata M. *et al* 2005 *Nucl. Fusion* **45** 1056
[421] Parks P.B. 2005 Gas jet pressure requirement for direct penetration, private communication (General Atomics, San Diego, USA)
[422] Kuteev B.V., Sergeev Yu V. and Sudo S. 1995 *Nucl. Fusion* **35** 1167