

EDGE LOCALIZED MODES IN TOKAMAKS

Howard Wilson

York Plasma Institute, Department of Physics, University of York, Heslington, York YO10 5DQ UK
Email: howard.wilson@york.ac.uk

ABSTRACT

As one increases the heating power in a tokamak beyond a threshold, the confinement undergoes a bifurcation, with a dramatic increase in the confinement time by a factor ~ 2 . This improved confinement regime, or H-mode, is primarily due to the formation of an insulating region at the plasma edge, where steep pressure gradients can form. A feature of H-mode operation is a series of explosive plasma eruptions, called Edge Localised Modes, or ELMs. They repeatedly expel large amounts of energy and particles from the plasma, with serious consequences for the heat loads that plasma facing components must be designed to handle. The present understanding of these ELMs in terms of ideal magneto-hydrodynamic instabilities will be reviewed in this paper.

I. INTRODUCTION

In the early 1980's an exciting new discovery was made in tokamaks: the H-mode¹. This mode of operation is achieved by increasing the plasma heating power above a threshold. The plasma then undergoes a bifurcation to a new state called the H-mode, in which the confinement is typically a factor of two greater than that in the standard L-mode discharges. The discovery of the H-mode is exceedingly important for fusion. In particular, ITER relies on achieving H-mode in order to meet its objectives.

There has been much activity to understand why the plasma suddenly jumps to this high confinement state. It is known, for example, that the reason for the high confinement is that a transport barrier forms at the plasma edge. A transport barrier is a region where heat and/or particle transport across the magnetic flux surfaces is suppressed. The pressure gradient is therefore typically high there. While this is good for confinement, instabilities called edge localized modes, or ELMs, can be triggered^{2,3,4} in this high pressure gradient region (which is also sometimes called the pedestal region). These ELMs are repetitive events, which eject a large amount of

heat and particles from the plasma. In a large tokamak, like ITER, the energy expelled could cause excessive erosion⁵. It is therefore essential that this phenomenon is understood so that steps can be taken to ameliorate the effect of ELMs, or perhaps eliminate them altogether. This paper describes the progress that has been made in our understanding of ELMs in terms of magneto-hydrodynamic (MHD) instabilities. In particular, we shall explore some of the characteristic features of ELMs, and how they can be understood in terms of this model. In Section 2, we discuss some of the experimental signatures for ELMs, before introducing the so-called peeling-ballooning model in Section 3. This linear ideal magneto-hydrodynamic (MHD) model is widely believed to explain at least the largest ELMs, and some of the experimental evidence for this is described in Section 4. In Section 5 we consider non-linear aspects of the model, and again compare with experimental observations. We conclude in Section 6, including a brief discussion on ELM control techniques.

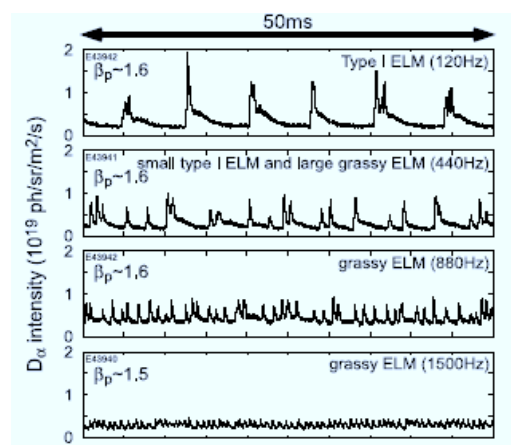


Fig. 1. D_α light emission, showing a sequence of ELMs on the JT-60U tokamak as bursts in the intensity; time traces for 4 discharges are shown giving examples of large, Type I ELMs (top) down to small, high frequency, “grassy” ELMs (bottom). [reproduced from Ref 6]

II. EXPERIMENTAL SIGNATURE OF ELMs

IIA. D_α light

An ELM is typically detected by observing the emission of D_α light near the divertor target plates, for example (see Fig 1). It is worthwhile considering the origin of the D_α light. The largest cross section for emission results from the interaction of electrons with neutral particles. Thus D_α light which is emitted from the divertor target plate region arises due to the electrons which are ejected from the core plasma during the ELM and then rapidly travel along the scrape-off layer towards the target plates where they interact with neutral particles. Ions are also ejected during the ELM, and these travel down to the target plate on a longer time-scale. They also increase the level of D_α light there by either (i) releasing more neutral particles from the target plate, which then interact with the electrons, or (ii) become neutralized themselves, again increasing the number of neutral particles for the electrons to interact with. After rapidly reaching a peak, the D_α light decays away more slowly as the ELM event finishes and the number of electrons arriving at the target plates decays to zero.




Fig. 2. Plot of temperature pedestal versus density pedestal for a range of different ELM types on DIII-D [figure reproduced from Ref 8]

IIB. ELM types

Experimentalists have managed to group the range of different ELM types into a number of different categories. Most ELMs fall into the category of Type I or Type III. The way these are typically distinguished is by measuring how the frequency of ELM events (ie the “ELM frequency”) varies with heating power. For Type I ELMs, the frequency increases with power, while for Type III ELMs the frequency decreases. Perhaps a more useful

distinction between the ELM types is obtained by comparing the temperatures and densities of discharges⁷. These are measured at the inner edge of the transport barrier: the so-called pedestal values. Figure 2 shows an example of how a set of discharges from DIII-D are distributed in the space of temperature pedestal, T_{ped} , versus density pedestal, n_{ped} ⁸. Note that the Type I ELMs are clustered around a line of constant pressure. This provides some evidence that these are instabilities that are driven by the plasma pressure. There are two clusters of Type III ELMs: one at high density, low temperature, and one at low density, high temperature. Both of these clusters are positioned at a pedestal pressure significantly below that where Type I ELMs occur. Note also that the Type III ELMs occur in a similar region to where the transition from the L-mode to the H-mode occurs.

It is found that regimes with Type I ELMs tend to have better confinement than those with Type III, but that the energy expelled in a Type I ELM event is much greater than for a Type III. Indeed, Type I ELMs would likely be too damaging on ITER, and therefore cannot be tolerated. The confinement degradation caused by Type III ELMs, on the other hand, is likely to be unacceptable. From Fig 2, it is clear that the pressure pedestal in Type III-ELMing discharges is significantly less than that in Type I discharges. This lower pedestal pressure is presumably the cause of the reduced confinement. This therefore raises the question “can we operate with a pressure pedestal characteristic of Type I ELM regimes, but with ELMs that release a much smaller amount of energy”? The answer is “yes”: but it is not easy and the parameter space required to achieve this is presently rather restrictive. Nevertheless, a range of small ELM regimes exist which seem to have good confinement⁹. These include Type II ELMs, grassy ELMs, Type V ELMs, etc. An example of grassy ELMs from the Japanese tokamak, JT-60U, is shown in the lower D_α trace of Fig 1, and compared to that for Type I ELMs. No clear definition of these regimes exists, and we will not attempt to refine our definition beyond that of small ELMs with good confinement. There is as yet no complete theoretical understanding for these ELM-types, but they do seem to be associated with strong plasma shaping and/or high edge safety factor. In addition, the grassy ELMs shown in Fig 1 appear to be associated with higher plasma rotation⁶. Without a rigorous understanding, it remains unclear whether ITER will be able to operate with any of these small ELMs. The subject therefore remains an active area of tokamak research.

III. PEELING-BALLOONING MODEL OF ELMs

Over the past decade the peeling-ballooning model^{10,11,12,13} has emerged as the strongest contender to explain the characteristics of the largest ELMs: Type I. In

addition, there is scope within the model to explain the origin of smaller ELM types, although the models here are much more qualitative and not universally accepted.

As the name suggests, the model derives from two particular instabilities: the ballooning mode and the peeling mode. The ballooning mode has a short wavelength perpendicular to the magnetic field lines, and a long wavelength parallel to them. It is destabilized when the pressure gradient exceeds a critical value, but tends to be stabilized by current density. In particular, sufficiently high current density completely stabilizes the ballooning mode, allowing the pressure gradient to be increased indefinitely (neglecting the effect of other types of instability). This is known as “second stability access”. These features are illustrated in Fig 3a. Actually, this second stability access strictly only occurs in the limit that the toroidal mode number, n , tends to infinity. As we shall see later, intermediate n modes (the peeling-ballooning modes) also have a drive due to the current density gradient, and these typically close off the second stability access. A final point to note is that the instability typically affects the whole of the plasma in the transport barrier region, and would therefore be expected to have a significant impact on it.

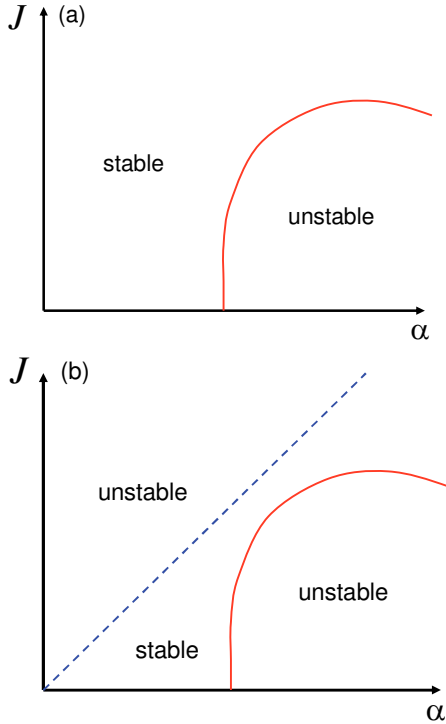


Fig. 3. Stability diagrams for (a) the ballooning mode and (b) the ballooning (full) and peeling mode (dashed). J is the current density at the plasma edge, α is the normalized pressure gradient.

The second type of instability, the peeling mode, is destabilized by the current density at the plasma edge, but stabilized by pressure gradient there. It is strongly related to the external kink instability, but is not limited to finite n . The schematic stability diagram for a limited plasma is illustrated in Fig 3b. This peeling mode is highly localized, and only affects a very small region of the transport barrier, exceedingly close to the plasma edge. At high pressure gradient, the peeling and ballooning modes couple^{10,14}, providing a strong instability with both current and pressure gradient drives that affects the whole transport barrier region. It is this coupled instability that is thought to be responsible for driving Type I ELMs.

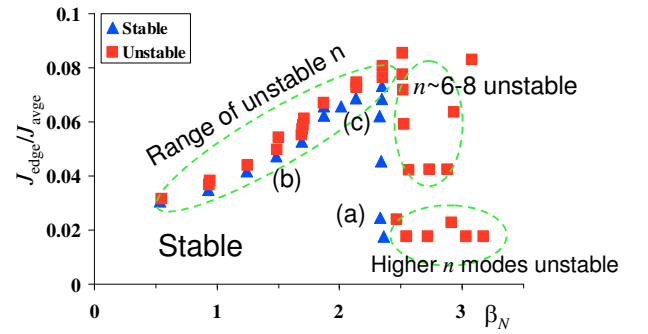


Fig. 4. Stability diagram for a JET-like plasma. The edge current density, J_{edge} , and normalized pressure, β_N , for each equilibrium are plotted as a square if unstable, and a triangle if stable.




Fig. 5. ELITE calculation for the mode structure in the poloidal plane for a coupled peeling-ballooning mode. Orange represents zero perturbation, while dark and light colours correspond to positive and negative perturbations, respectively.

In Fig 4 we show the stability diagram for a (JET-like) model tokamak equilibrium. This has been produced using the ELITE code^{12,15}, which is a very efficient code for calculating the stability of the plasma edge to ideal MHD instabilities with intermediate to high toroidal mode numbers, n . For each point, the current density in the pedestal region is assumed to be a combination of the bootstrap current and Ohmic current. The stability boundary (between square and triangle symbols) is qualitatively similar to what one would expect, combining the peeling and ballooning stability boundaries of Figs 3b. Using ELITE, we can also evaluate the mode structures. Figure 5 shows the mode structure for a peeling-ballooning mode, which extends right across the pedestal region. Note that the amplitude is maximum on the outboard (low field) side. This characteristic is referred to as “ballooning”. The pure ballooning mode has a similar structure, but the pure peeling mode differs in two respects. First, the peeling mode has a similar amplitude on the inboard side to that on the outboard side (ie no ballooning nature). Second, it is extremely edge localized, typically well within the last 1% of poloidal flux (unless the toroidal mode number is low, ie $n \sim 1, 2$); we do not plot it here as it would not even be visible. The pure ballooning modes, from the region of Fig 4 labeled (a), have very high toroidal mode numbers (several 10’s); the pure peeling modes, from the region labeled (b) have a range of toroidal mode numbers, the most unstable having their external rational surface closest to the plasma edge; the peeling-ballooning modes, from the region labeled (c), have intermediate toroidal mode number $n \sim 10$.




Fig. 6. Possible trajectories (dashed curves) of edge current density, J , and normalized pressure gradient, α , during ELMs. The stability boundary is indicated by the full curve. See text for descriptions and suggested links to ELM types.

Stability diagrams such as Fig. 4 are a simplification of the actual situation, as the stability diagram typically depends on the width of the pedestal region, as well as the pressure gradient within it (narrower pedestals tend to be more stable than wider ones at fixed pressure gradient).

The linear stability calculations shown in Figs 4 and 5 motivate the peeling-ballooning model for ELMs. Figure 6 shows a number of possible trajectories for the edge pressure gradient and current density in an ELMing discharge. These could correspond to different ELM types as follows. Starting with the trajectory labeled 1 in Fig 6, this is proposed as the trajectory of edge plasma parameters during a Type I ELM. Thus, while the plasma is stable between ELMs, the edge pressure gradient increases up to the ballooning boundary. The bootstrap current, which is proportional to the pressure gradient, would also increase until the peeling-ballooning mode is destabilized. Recall that this mode extends right across the transport barrier, so affects the whole pedestal region. As the instability develops, the expected increase of transport causes the pressure gradient to fall, further destabilizing the mode and triggering a large crash in the pressure (and consequent large energy loss). The discharge parameters eventually re-enter the stable region, and the cycle repeats. The trajectory labeled 2 is somewhat different. The plasma achieves a similar pressure gradient (and hence confinement) to the trajectory labeled 1, but now only the ballooning mode is destabilized. In this regime, which has reduced current compared to trajectory 1, the drop in pressure gradient stabilizes the ballooning mode, switching off the instability without a dramatic loss in pressure. This could provide an explanation of some of the small ELM regimes. Finally, the trajectory labeled 3 could provide a substantial loss of pressure gradient but, because the instability is so extremely edge-localized, the total loss of energy might be small (no rigorous theory for this speculative statement exists as yet). Note that the pressure gradient for this trajectory is significantly below that for trajectory 1, and this might therefore provide a model for Type III ELMs where the confinement is reduced (at least the class with low density and high temperature, where the edge current density is expected to be high).

Of the three models, the trajectory describing the Type I ELM behavior is the most accepted. The others are more speculative. Indeed, we shall see later that there is additional important physics that has been deduced from non-linear theories. This should be taken into account for a full understanding of ELM behavior. Before we address this non-linear physics, let us first consider some of the experimental evidence that supports some of the features predicted by this linear MHD model for Type I ELMs.

IV. EXPERIMENTAL EVIDENCE FOR THE PEELING-BALLOONING MODEL

The best way to test the peeling-ballooning model for ELMs is to carefully reconstruct the equilibrium using detailed measurements of the current density and pressure

gradient in the transport barrier, and compare with the predicted stability boundaries. The pressure increases by an order of magnitude over the ~few cm width of the transport barrier so that high resolution, high accuracy measurements of the density and temperature are required. Note that the error in the pressure gradient is significantly larger than the error in the pressure. Nevertheless, there are several tokamak experiments around the world that can now make this measurement with a reasonable accuracy (though not always for both the electrons and ions, so that often some form of approximation must be made). The other key requirement is the current density. In the core plasma, this is usually inferred from a measurement of the magnetic field using a technique such as motional Stark effect (MSE). This type of measurement is significantly complicated in the transport barrier for two reasons. First, the magnetic field structure in this region is determined largely by the geometry produced by the magnetic coils, and is less sensitive to the current density in the plasma. Second, there is a strong electric field in the transport barrier (indeed, this is widely thought to be responsible for the formation of the transport barrier). This electric field must be taken into account when interpreting the Stark splitting of the spectra. On DIII-D significant progress has been made in the measurement of current density using a lithium beam¹⁶, but this is not a routine measurement, and interpretation of the data is difficult. Generally, then, the current density in the transport barrier is not known, and one must resort to modeling.




Fig. 7. The D_α trace of a DIII-D discharge and the ideal MHD growth rate (full curve) from an ELITE stability analysis [12]

The bootstrap current is driven in a tokamak through a combination of the trapped particles and the pressure gradient. Although it is suppressed to some extent in the transport barrier through collisional effects, it is nevertheless typically the largest contribution because of

the high pressure gradient that exists there. Therefore most stability analyses of the plasma edge region simply use the bootstrap current expression rather than a direct measurement of the current density.

We consider two examples of how the peeling-ballooning model is typically tested against experimental data. The first is a comparison with data from DIII-D. The upper trace in Fig 7 shows the D_α light as a function of time. The drop in the signal at the beginning indicates the transition from L mode to H-mode, at which time the steep pressure gradient begins to form in the transport barrier region. A sequence of stability analyses using the ELITE code (setting $n=10$, which is typically the most unstable mode number) indicates that the plasma is stable to peeling-ballooning modes up to the time 2150ms, at which point the plasma becomes unstable, and the growth rate begins to rise until the first ELM appears at a time of around 2300ms. Thus, there is a reasonable correspondence between the predicted onset of instability and the onset of the ELM (particularly when diamagnetic effects are taken into account, in which case the growth rate must exceed half of the diamagnetic frequency, $\omega_*/2$, for instability).

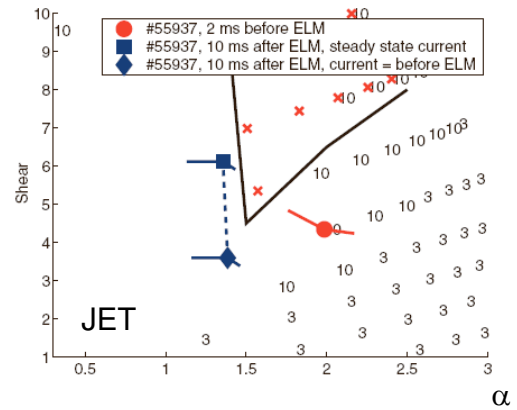


Fig. 8. Stability diagram for a JET discharge. See text for description. [Reproduced from ref 17]

The second type of stability analysis that we show is more typical of the analyses that have been performed on tokamaks like JET¹⁷, DIII-D¹², MAST¹⁸ and ASDEX-Upgrade¹⁹, for example. The specific case shown in Fig 8 corresponds to a study on JET¹⁷. The diagram shows the magnetic shear (which increases for decreasing current density) and normalised pressure gradient, α , for a range of equilibria. Each unstable equilibrium is represented by a single number, which indicates the most unstable toroidal mode number. The solid curve denotes the $n=\infty$ ballooning mode stability boundary. The areas with no numbers, not bound by the solid curve are stable. Also shown on the figure are three points representing the

actual experimental values of shear and α for this JET discharge (55937). The circle shows the parameters immediately (2ms) before an ELM, indicating that the discharge is indeed close to the stability boundary. It is therefore reasonable to assume that the instability is the cause of the ELM. Just after the ELM event (10ms after), the pressure gradient is lower, and the discharge moves away from the stability boundary. There are two points indicated at this later time, and these correspond to two assumptions about the plasma current which, as stated earlier, cannot be measured. The point at higher magnetic shear has a lower plasma current, given by the steady bootstrap current associated with the lower pressure gradient at this later time. The point at lower magnetic shear has a higher current density: in this case it is assumed that in the brief period between the ELM onset and the end of the ELM event, the current has not had time to relax. The current is therefore assumed to be the same as that immediately before the ELM (and given by the steady state bootstrap current at that time).

V. NON-LINEAR MODELS

VA. The theory

While linear models provide quantitative information about the onset of instability, they tell us relatively little about the consequences. In particular, it is unlikely that we shall be able to construct a predictive model for the energy ejected during an ELM on ITER based purely on linear theory. Having said that, we have already suggested that there might be a link between the radial extent of the linear eigenmode structure and the energy lost during the ELM. There is certainly some evidence to support this, at least qualitatively^{20,21}, but also quantitatively in some cases (eg from DIII-D¹²). However, in general there does not seem to be a one-to-one correspondence between the radial width of the linear eigenmode and the ELM-affected volume²². This does not exclude some form of correlation between the radial eigenmode width and the ELM-affected region, but the two do not appear to be the same in general.

Thus it seems unlikely that we can address the key question of how much energy is ejected during an ELM from a model based purely on the linear theory. This has motivated recent studies that address non-linear models. One of the first was an analytic calculation during the early non-linear phase of the ballooning mode in a tokamak plasma²³. This has recently been extended deeper into the non-linear regime²⁴. There were two key predictions from this first model²³: that the ballooning mode grows explosively during the non-linear phase, and that the instability would eject hot filaments of plasma. The explosive nature is an important point. One might expect from linear theory that a pure ballooning mode

would grow relatively slowly as the pressure gradient builds on a transport timescale through the marginally stable value. However, the non-linear theory suggests there is much more to the story. As the mode amplitude grows to a level where the non-linear terms first become important, the growth is accelerated much above the linear growth rate. In addition, the spatial structure of the mode is also altered during the non-linear phase to form filament-like structures that are predicted to be aligned with the magnetic field lines.

In Section 3 we described a model based on linear theory which suggests that the crash is a consequence of a coupling between the peeling and ballooning modes. The explosive growth predicted by the non-linear model suggests that it is possible to trigger a crash event just by considering the ballooning mode. However, there is another twist. The equation that describes the non-linear evolution of the ballooning mode has a key term that causes the explosive behavior. If the coefficient of this term is positive, the filaments erupt outwards towards the scrape-off layer, otherwise they erupt inwards towards the core. It has been found theoretically that filaments only erupt outwards when there is sufficient current density in the plasma edge²⁵. If one assumes that an outward-erupting filament is more dangerous than an inward one, then this could provide another reason why the ELM size might depend on the edge current density and provide an ingredient to understanding the differences between the mechanisms for large and small ELMs.

There are other types of models that predict filamentary structures. These are related to so-called “blob” theories, and predict that if a filament of plasma should break off from the core plasma, then it will propagate radially outwards due to an ExB drift^{26,27}. This is different to the non-linear ballooning mechanism, where the filament is only ejected from the outboard side, and remains connected into the core plasma on the inboard side. Thus the filaments due to a non-linear ballooning mode can continue to tap the free energy of the pressure gradient in the transport barrier, and accelerate rapidly away from the core plasma. In addition, because they remain connected to the core plasma, they can act as a conduit (or hose-pipe), linking the transport barrier region to the scrape-off layer, siphoning hot plasma from the barrier region into the exhaust region. If this is the mechanism, it raises a new concern: what if the filaments should strike the vessel wall on the outboard side while still connected to the core plasma on the inboard side? This would place very high heat loads on the vessel walls, where there is relatively little protection. If, on the other hand, the filaments break off from the core plasma before striking the wall they would have a reduced impact on it. If this is the case, there is scope for both mechanisms to play a role: the non-linear ballooning theory may be

relevant for the early phase of the ELM, while the “blob” theories may be relevant later, after the filament has detached.




Fig. 9. Mirnov signals from a series of coils positioned at various toroidal and poloidal positions around the COMPASS-D tokamak²⁸.

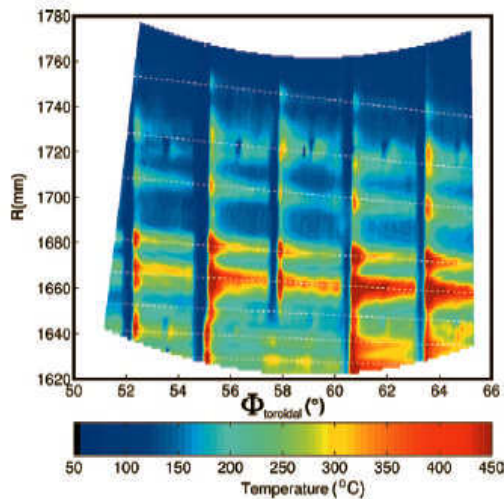


Fig. 10. Filamentary structures evident in the infra-red emission from the outer sections of the ASDEX-Upgrade divertor target [Figure reproduced from Ref 30]

VB. Experimental evidence for filaments

Some of the first data that suggested that filamentary structures might be associated with ELMs comes from COMPASS-D²⁸. A poloidal and toroidal array of Mirnov coils allowed the magnetic signal associated with the ELM to be tracked in time in the two directions. The results, reproduced in Fig 9, indicate a field-aligned structure. More recently, a series of stripes have been observed in the infra-red emission from the outer sections of the divertor in ASDEX-Upgrade (see Fig 10)^{29,30}. These stripes are also consistent with ELMs ejecting field-aligned filamentary structures. The most dramatic evidence, where the filaments were first observed directly, comes from the MAST tokamak^{31,32,33}. A unique feature of a spherical tokamak is that one can view the whole plasma surface through a port. Such a view is shown in Fig 11 where the filamentary structures are clearly visible. These filaments are aligned with the magnetic field lines, as theory would predict³³, and are ejected far into the scrape-off layer, and beyond. They are observed to accelerate radially, qualitatively consistent with the prediction of non-linear ballooning theory, and decelerate in the toroidal direction³³.




Fig. 11. Filamentary structures observed during ELMs on the MAST tokamak³²

Another interesting observation, made on both MAST and DIII-D, is that the rotation shear (that is the differential rotation between different flux surfaces) is eliminated during the ELM event^{34,35} (see Fig 12). Within the ideal MHD model of the non-linear ballooning mode, this must happen: the filament could not erupt out of the plasma if there is differential rotation. So again there is a nice consistency here. It does, however, raise an interesting question: Is the eruption of the filament causing a locking of the flow profile in the outer regions of the plasma, or is there some other event which removes

the flow shear, which then allows the filaments to erupt? In other words, is the filament a cause or consequence of the ELM? The consensus is that the filaments are the cause, but conclusive proof remains elusive.

VC. Heat loss mechanisms

Let us close this section with a discussion of the possible heat loss mechanisms. We simply state the possible mechanisms which might be operating here, and do not attempt to rule any of the models out or in. More details are provided in Ref 36.




Fig. 12. Flow shear suppression during ELMs on the MAST tokamak. The full circles are the earliest time indicated by the vertical lines in the D_α trace of a single ELM (top), followed by the full squares, full triangles, open squares and open triangles. [Fig reproduced from Ref 34]

A key point to note is that within ideal MHD there can be no loss of heat or particles: everything that flows out along the filament into the scrape-off layer would flow back into the confinement region while the filament remained attached to the core plasma on the inboard side. It is necessary to introduce some additional physics. Three mechanisms are proposed for how the filaments might lead to enhanced transport. (1) The hot filaments “leak”, so that hot plasma flows across field lines from

the filament as it pushes out into the scrape-layer. (2) There is a reconnection process, probably in the vicinity of the X-point, in which the magnetic field lines contained in the hot filament break and then reconnect with the magnetic field lines of the scrape-off layer on the outboard side. Because the filament is relatively unperturbed on the inboard side, it remains inside the transport barrier there. Thus, following the reconnection event, there would be a continuous path along the filament from the pedestal region to the divertor target plates, and therefore a rapid loss of heat and particles. (3) The transport barrier is thought to be sustained by a sheared plasma flow within it. As the filament pushes out, it suppresses the flow shear and the barrier collapses with a significant, temporary confinement degradation. There is much experimental activity, and increasing efforts to simulate ELMs computationally, to identify the dominant mechanism. Both are exceedingly difficult, and are high priority areas of research for ITER.

VI. CONCLUSIONS

The importance of understanding ELMs and developing control scenarios for ITER is now well-established. Considerable progress has been made in our understanding, with the majority of the community accepting the peeling-ballooning model for at least the largest, Type I ELMs. There is also scope within the model to explain some of the other ELM regimes, but there is more uncertainty here. Non-linear theories have been developed and, together with experiment, have established that filamentary structures are likely to be an important part of the process of energy and particle loss during ELM events.

So, after all of this work, what are the consequences for ITER? ELMs remain one of the biggest concerns for ITER, which cannot operate for more than a few discharges at full power if large Type I ELMs are triggered. Small ELM regimes do exist, but the operating space for these is generally small, and it is not even clear that they will extrapolate to ITER. For this reason, there has been an increased focus on ELM control. We do not go into the details here, but instead refer the interested reader to the references. One method involves applying small magnetic perturbations to the plasma edge, thus degrading the confinement. The mechanism is not completely understood, but ELMs are then generally suppressed, probably because the pressure gradient in the transport barrier is maintained below the critical value for instability. Experiments on DIII-D validated and then perfected this form of ELM control^{37,38}, which has recently been repeated in ASDEX-Upgrade³⁹. A second technique, pioneered on ASDEX-Upgrade, is complementary to this. Rather than suppress ELMs, they are instead triggered more frequently by firing small

pellets of frozen deuterium into the transport barrier. Each time a pellet is fired into the plasma edge, an ELM is triggered⁴⁰. By triggering frequent ELMs, the energy released in a single event is smaller, and the impact on the divertor target is much reduced. A physical understanding of the mechanism is emerging from nonlinear MHD simulations⁴¹.

To summarize, there is still much to do for a complete understanding of the ELM event, although much progress has been made in the last decade. As well as being an issue that we must address for ITER, it is an area that is rich in physics. Here we have focused on the plasma physics issues, but there is also a range of issues related to plasma-surface interactions that must be brought into the model to develop a complete understanding of the ELM cycle: a formidable, but rewarding, task.

ACKNOWLEDGMENTS

I am grateful to Andrew Kirk, for many enlightening discussions regarding the experimental data and its interpretation.

REFERENCES

1. F. WAGNER, et al, *Phys Rev Lett* **53**, 1453 (1984)
2. H. ZOHM *Plasma. Phys. Control. Fusion* **38**, 105 (1996)
3. J. W. CONNOR, *Plasma. Phys. Control. Fusion* **40**, 531 (1998)
4. W. SUTTROP, *Plasma. Phys. Control. Fusion* **42**, A1 (2000)
5. A. LOARTE, et al, *Phys Scri* **T128**, 222 (2007)
6. N. OYAMA, et al, *Nucl Fusion* **45**, 871 (2005)
7. W. SUTTROP, et al *Plasma Phys. Control. Fusion* **39**, 2051 (1997)
8. T. OSBORNE, et al, *Proceedings of the 24th EPS Conference on Controlled Fusion and Plasma Physics, Berchtesgaden* Part III, 1101 (1997)
9. N. OYAMA, et al, *Plasma Phys. Control. Fusion* **48**, (2006) A171
10. J.W. CONNOR, et al, *Phys. Plasmas* **5**, 2687 (1998)
11. H.R. WILSON, et al, *Phys. Plasmas* **6**, 1925 (1999)
12. P.B. SNYDER, et al, *Phys. Plasmas* **9**, 2037 (2002); Figure reprinted with permission from reference. Copyright (2002), American Institute of Physics.
13. P.B. SNYDER, et al, *Nucl Fusion* **44**, 320 (2004)
14. C.C. HEGNA, et al, *Phys Plasmas* **3**, 584 (1996)
15. H.R. WILSON et al, *Phys. Plasmas* **9**, 1277 (2002)
16. D. THOMAS, et al *Phys. Rev. Lett.* **93**, 065003 (2004)
17. S. SAARELMA, et al, *Plasma Phys. Control. Fusion* **47**, 713 (2005)
18. S. SAARELMA et al, *Plasma Phys. Control. Fusion* **49**, 31 (2007)
19. C. KONZ, et al, to appear in *Proceedings of the 34th EPS Conference on Controlled Fusion and Plasma Physics, Warsaw* (2007)
20. L.L. LAO, et al, *Nucl. Fusion* **41**, 295 (2001)
21. S. SAARELMA et al, *Nucl. Fusion* **43**, 262 (2003)
22. A. LOARTE, et al *Plasma Phys. Control. Fusion* **44**, 1815 (2002)
23. H.R. WILSON and S.C. COWLEY, *Phys Rev Letts* **92** 175006 (2004)
24. P. ZHU, C.C. HEGNA and C.R. SOVINEC, *Phys Plasmas* **14**, 055903 (2007)
25. H.R. WILSON, et al *Proceedings of the 21st IAEA Fusion Energy Conference, Chengdu, on CD-ROM: IAEA-CN-149*, paperTH/4-1Rb (2006)
26. S.I. KRASHANINNIKOV, *Phys Lett A* **283**, 368 (2001)
27. W. FUNDAMENSKI, et al *Plasma. Phys. Control. Fusion* **48**, 109 (2006)
28. M. VALOVIC, et al, *Proceedings of 21st EPS Conference on Controlled Fusion and Plasma Physics, Montpellier* Part I, 318 (1994)
29. T. EICH, et al, *Phys Rev Lett* **91**, 195003 (2003)
30. T. EICH, et al, *Plasma. Phys. Control. Fusion* **47**, 815 (2005)
31. A. KIRK, et al, *Phys Rev Lett* **92**, 245002 (2004)
32. A. KIRK, et al *Proceedings of the 21st IAEA Fusion Energy Conference, Chengdu, on CD-ROM: IAEA-CN-149*, paper EX/9-1 (2006)
33. A. KIRK, et al *Plasma. Phys. Control. Fusion* **48**, B433 (2006)
34. A. KIRK, et al *Plasma. Phys. Control. Fusion* **47**, 315 (2005)
35. J.A. BOEDO, et al, *Phys. Plasmas* **12**, 072516 (2005)
36. H.R. WILSON, et al, *Plasma Phys. Control. Fusion* **48** A71 (2006)
37. R.A. MOYER, et al, *Phys. Plasmas* **12** (2005) 056119
38. T. E. EVANS, R.A. Moyer et al. *Phys. Rev. Lett.* **92** (2004) 235003
39. W. SUTTROP et al, *Phys. Rev. Lett.* **106** (2011) 225004
40. P.T. LANG, et al, *Nucl. Fusion* **44** 665 (2004)
41. G.T.A. HUYSMANS, et al, *Plasma Phys. Control. Fusion* **51** (2009) 124012