Resonant-like behaviour during edge-localised mode cycles in the Joint European Torus

A. J. Webster, J. Morris, T. N. Todd, S. Brezinsek, P. Coad, J. Likonen, M. Rubel, and JET EFDA Contributors

Resonant-like behaviour during edge-localised mode cycles in the Joint European Torus

A. J. Webster,[1] J. Morris,[1] T. N. Todd,[1] S. Brezinsek,[2] P. Coad,[1] J. Likonen,[3] M. Rubel,[4] and JET EFDA Contributors[a]

JET-EFDA, Culham Science Centre, Abingdon OX14 3DB, United Kingdom
[1]*CCFE, Culham Science Centre, Abingdon OX14 3DB, United Kingdom*
[2]*Institut für Energie- und Klimaforschung - Plasmaphysik, Forschungszentrum Jülich, 52425 Jülich, Germany*
[3]*VTT, Association Euratom-Tekes, PO Box 1000, FI-02044 VTT, Finland*
[4]*Alfvén Laboratory, School of Electrical Engineering, Royal Institute of Technology (KTH),*
Association EURATOM-VR, Stockholm, Sweden

A unique sequence of 120 almost identical plasmas in the Joint European Torus (JET) recently provided two orders of magnitude more statistically equivalent data than ever previously available. The purpose was to study movement of eroded plasma-facing material from JET's new Beryllium wall, but it has allowed the statistical detection of otherwise unobservable phenomenon. This includes a sequence of resonant-like waiting times between edge-localised plasma instabilities (ELMs), instabilities that must be mitigated or avoided in large magnetically confined plasmas such as those planned for ITER. Here, we investigate the cause of this phenomenon, using the unprecedented quantity of data to produce a detailed picture of the plasma's behaviour. After combining the data, oscillations are clearly observable in the plasma's vertical position, in edge losses of ions, and in Beryllium II (527 nm) light emissions. The oscillations are unexpected, are not obvious in data from a single pulse alone, and are all clearly correlated with each other. They are likely to be caused by a small vertical oscillation that the plasma control system is not reacting to prevent, but a more complex explanation is possible. The clearly observable but unexpected link between small changes in the plasma's position and changes to edge-plasma transport and stability suggest that these characteristics cannot always be studied in isolation. It also suggests new opportunities for ELM mitigation and control that may exist. [http://dx.doi.org/10.1063/1.4927755]

I. INTRODUCTION

In July 2012, more than 120 2T, 2MA H-mode plasmas were consecutively created over a 2-week period in the Joint European Torus (JET). The plasmas were designed to study the movement of eroded plasma facing material within JET's vacuum vessel[1] after the installation of JET's new ITER-like wall,[2] and each had about 6 s of H-mode of which 2–3 s at the end was exceptionally steady. Combining the data from the 2 s period prior to the end of H-mode from each of the 120 pulses provides 240 s of plasma with ∼10 000 edge-localised plasma instabilities (ELMs). In contrast, most plasma experiments produce only 2–3 identical pulses, sometimes from different sessions on different days or weeks. By synchronising data to the ELM times and averaging the result, the central limit theorem ensures that random errors are reduced by a factor of $1/\sqrt{10000} = 0.01$, if 10 000 pairs of ELMs are studied. Consequently, the large quantity of ELM data enables an unprecedentedly detailed description of the average plasma behaviour prior to and after an ELM. Full details of the plasmas are in Refs. 1 and 3.

An unexpected result[3] was the sequence of sharp maxima and minima in the probability density function (pdf) for the waiting times between ELMs, which was formed by binning the waiting-times between ELMs into 0.005 s intervals

and normalising the resulting distribution (Figure 1). With only 150–200 ELMs, as are typically produced in 2–3 type-I JET H-mode plasmas, a pdf with the resolution of Figure 1 would have large error bars that would prevent us from distinguishing "structure" from statistical noise. This presently prevents similar accurate studies of JET plasmas due to an insufficient: number of ELMs, number of repetitions of identical plasmas, or length of steady-state plasma phase. For
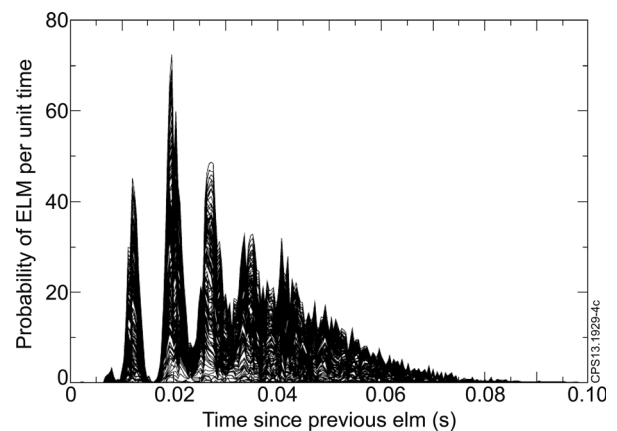

FIG. 1. The ELM waiting-time data from 120 almost identical JET plasmas is combined to form a single pdf for the waiting times between ELMs. Each line corresponds to data from an additional pulse, which are added together to form the pdf. Reproduced with permission from Webster *et al.*, Plasma Phys. Controlled Fusion **56**, 075017 (2014). Copyright 2014 EURATOM.

example, previous work[4] examined 85 examples of high quality JET H-mode plasma pulses and found that the waiting time data could be well fitted with a Weibull distribution. Because the number of type-I ELMs in each pulse was only of order 50–100, if a finer structure was present in the pdfs, the only observable symptom would be a poorer fit between the data and the assumed (Weibull) pdf.

The phenomenon becomes clearer as the plasma becomes increasingly steady-state, as can be seen by plotting the ELM waiting times versus time through the plasma pulse for all 120 plasmas (Figure 2). The horizontally extended clusters corresponding to maxima and minima in Figure 1 become clearer after about 3 s. As mentioned previously, we only use the last ∼2 s of data from each pulse. The ELMs' properties are discussed in Ref. 5. One surprising observation was that for waiting times greater than 20 ms the losses of energy and density due to an ELM were found to be independent of their waiting time. The energies of subsequent ELMs were also found to be independent of each other.

Here, the large amount of (statistically equivalent) ELM data is used to provide a detailed picture of the plasma's behaviour prior to an ELM, with the hope of better understanding the causes of the resonant-like maxima and minima in Figure 1. This is done by identifying ELMs from Be II light observed at the inner divertor with the method described in Ref. 4, then using the ELM times to synchronise the statistically equivalent data to form an average of the result. For each plot, we show the average from across the set of pulses, the standard deviation about it, and an example of data from the last pulse of the set (83 794). The pulse 83 794 was chosen because it was convenient to include, its behaviour would be expected to naturally differ slightly from the average. Details of how the plots are formed are in Appendix A, an explanation for the statistical accuracy of the method is in Appendix B, and a discussion of the selection of diagnostics and sources of uncertainty is in Appendix C.

Figure 3 plots the Be II light emissions synchronised to the ELM times. The plot shows a sequence of maxima and minima, with minima at intervals of approximately 8, 16, 24, and 32 ms, the same intervals as are found between minima in the ELM waiting time pdf in Ref. 3. Similar observations are found in Figure 4, which show 8 ms-period oscillatory losses in edge density that coincide with oscillatory fluxes of

FIG. 3. Beryllium light emissions measured at the inner divertor, synchronised to the ELM time ($t = 0$). The thick black line shows the average, the dashed black lines give the average plus/minus the standard deviation. The brown circles are points from a single typical pulse 83 794, which are plotted to give an example of how much an individual pulse deviates from the average. ELMs for which $t_n - t_{n-1}$ is less than 0.04 s are not included, to minimise the influence of the large post-ELM Be II response on the form of the pre-ELM signal ($t < 0$). Units are photons per second per cm^2 (vertical axis), and time in seconds (horizontal axis).
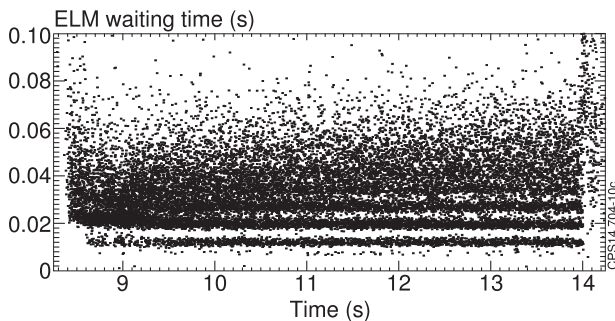

FIG. 2. The ELM waiting times (vertical axis) versus time through the pulse (horizontal axis), for all 120 pulses in the dataset. The vertically clustered horizontal stripes that correspond to the maxima and minima in Figure 1, become clearest after about 3 s into the (H-mode) pulse.
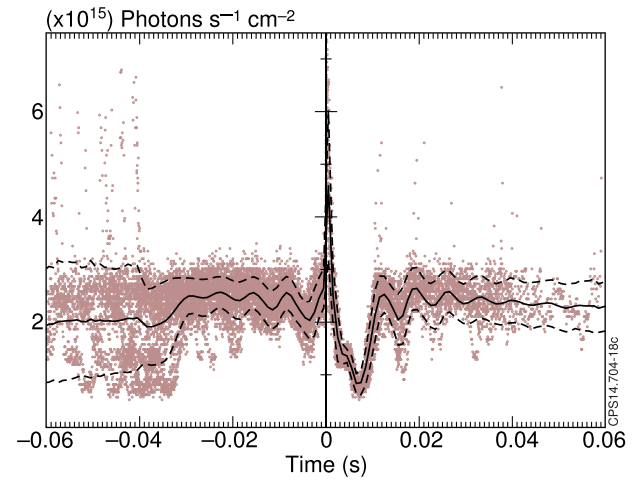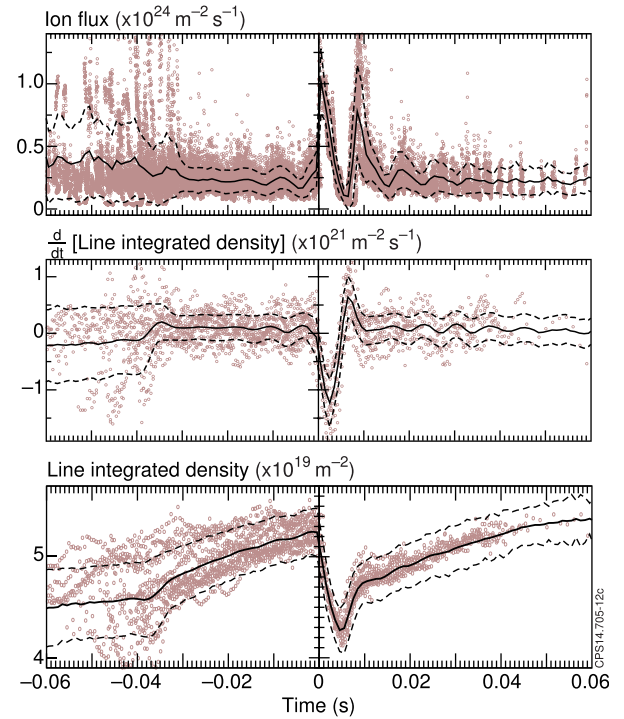
FIG. 4. Top—Langmuir probe measurements of ion flux to the inner divertor ($\times 10^{24}$ m^{-2} s^{-1}) versus time (s), synchronised to the ELM times at $t = 0$. Middle—the rate of change of line-integrated density measured through the plasma's mid-plane (particles m^{-2} s^{-1}) versus time (s), synchronised to the ELM times at $t = 0$. Bottom—the line-integrated density through the mid-plane (particles m^{-2}), versus time (s), synchronised to the ELM times at $t = 0$. Thick black lines are averages, dashed lines indicate standard deviations, and circles are typical measurements (pulse 83 794). Prior to ELMs there are 8 ms-period oscillations, with increased (or reduced) ion fluxes that coincide with reducing (or increasing) plasma density, similar to 8 ms-period plasma-position oscillations discussed later. The post-ELM signal is difficult to interpret due to large post-ELM plasma movements and non-linear interactions such as impurity influxes.

ions to the divertor, and with the pre-ELM Be II emissions shown in Figure 3. The observations could be explained if the plasma's stability is becoming increased (or reduced) at time intervals of approximately (8), 12, (16), 20, (24), 28, (32), 36 ms after an ELM, sometimes triggering an ELM, or otherwise leading to enhanced ion losses and Be II light emissions. It is known that ELMs can be destabilised by vertical displacements of the plasma,[6–9] offering one possible mechanism for modifying plasma stability. The rest of this article explores the evidence for a vertical plasma oscillation, and its association with the observed enhancement (or reduction) in ELM occurrence, Be II light emissions, and edge losses of ions.

II. EFIT EQUILIBRIUM RECONSTRUCTION

The plasma equilibrium shape is reconstructed from fast magnetic signals using EFIT,[10,11] and used to infer various quantities such as the gaps between the plasma and the vessel wall. Figure 5 shows changes to the inner, outer, and upper gaps between the plasma and vacuum vessel (offset to be zero at $t = 0$), and the vertical position of the X-point. An increase to the inner gap and a decrease to the outer gap correspond to a radial motion of the plasma outwards, and vice versa. An increase in the upper gap and a decrease in the vertical position of the X-point correspond to a downwards motion of the plasma, and vice versa. The plots in Figure 5 indicate that between ELMs there is on average a drift of the plasma outwards and downward (with associated increases in the inner and upper gaps, and decreases to the outer gap and the vertical position of the X-point). This is illustrated in Figure 6. Note that some of the distances we are discussing
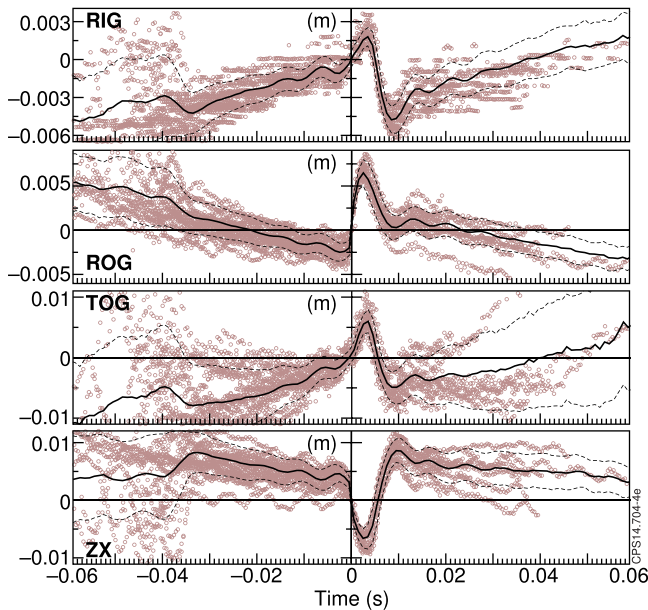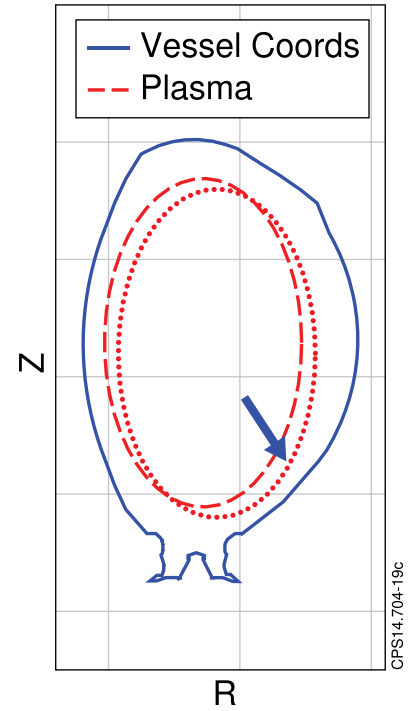

FIG. 6. Pictorial diagram of the diagonal plasma drift between the end of one ELM and the occurrence of another (not to scale). There is a 4–5 mm outward radial motion and a 4–7 mm downward vertical motion, resulting in a diagonal displacement by 5–8 mm. (The plasma never touches the wall.) The approximate displacements are from Figure 5.

are of order 1 mm in amplitude, an amplitude that is only observable because of the central limit theorem and the 1000 s of ELMs that were available for the analysis, as discussed earlier. In principle, the accuracy of EFIT can be improved by running with greater accuracy, including kinetic effects, and combining the outputs from other diagnostics to use these as constraints. Unfortunately, it is not presently possible to do this automatically on large numbers of pulses, preventing the statistical gains in accuracy that enable small oscillations to be observed. Fortunately, there are independent diagnostics that support EFIT's calculations, the most persuasive of which are presented in Section III. Other sources of uncertainty for EFIT are discussed in Appendix C.

In addition to the slow plasma drift radially outwards and vertically downwards, there are clearly visible oscillations with an approximately 8 ms period that correspond to a vertical motion downwards and a horizontal narrowing of the plasma. The vertical motion is illustrated in Figure 7. The radial narrowing and vertical displacements are a similar order of magnitude, and consequently, the measured increases in the radial inner and outer gaps (radial narrowing) do not seem to be due to a combination of vertical motion and plasma geometry (the plasma is radially narrower above and below the mid-plane). Because we do not include signals for which the times t_n and t_{n-1} of the nth and $(n-1)$th ELMs are separated by less than 0.04 s, the oscillations in the signal prior to an ELM are not due to the large displacements associated with ELMs. The physical cause of the oscillations is not certain, but their 8 ms period coincides with the period between the increased (decreased) likelihood of observing ELMs.

FIG. 5. Edge plasma displacements, offset so that they are zero at the ELM time ($t = 0$). From top to bottom: the radial inner gap (RIG), radial outer gap (ROG), top gap (TOG), and vertical X-point position (ZX). Circles are from a typical pulse (83 794), the solid line is an average across the pulse set, and the dashed lines are the average ± the standard deviation. The distances (vertical axis) are measured in metres, and the time (horizontal axis) is measured in seconds.
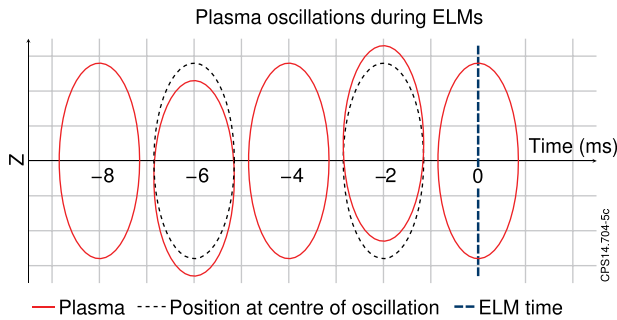
FIG. 7. Pictorial diagram to show the phase and timings of the vertical component of plasma oscillations that are observed before ELMs (not to scale). A single 8 ms cycle of the oscillation is shown. The timings are taken from plot 4 of Figure 5, which is the vertical X-point position. The plasma is moving with the oscillation's maximum downward velocity just before the ELM occurs (see plots 2 and 4 of Figure 10).

Figure 8 plots the radial and vertical position and velocity of the centre of plasma's current, as calculated by EFIT. It shows a radial drift outwards and downwards between ELMs, as was observed from changes to the position of the plasma's edge. There is no evidence of systematic radial oscillations, but there are clearly visible vertical oscillations, consistent with a vertical motion of the plasma. No conclusive evidence has yet been found for the vertical control system[12] reacting to the oscillations discussed here.[3] Figure 9 plots the signal used by the vertical control system as a proxy for the plasma's vertical velocity, and finds no obvious sign of any vertical oscillations, which would explain why there is no clear evidence of the vertical control system reacting to them. However, the velocity is estimated in a very different way to EFIT's calculation and uses a smaller subset of signals to do so. The signal used by the vertical control system
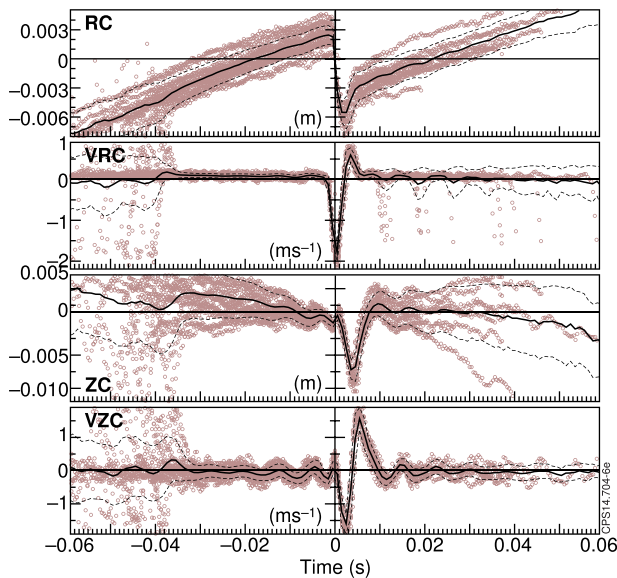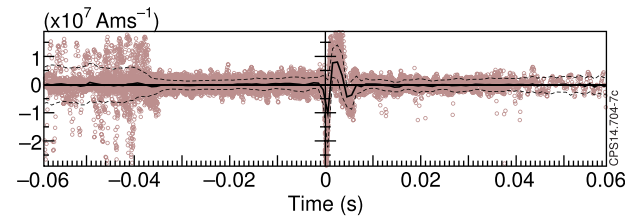

FIG. 9. The vertical control signal uses an approximation for the plasma's velocity, by estimating the rate of change of the product of plasma current and vertical position. No obvious vertical oscillations are observed by it. Circles are from a typical pulse (83 794), the solid line is an average across the pulse set, the dashed lines are the average $\pm$ the standard deviation, and the signals are offset so that they are zero at the ELM time ($t = 0$). The units are A m s^{-1} (vertical axis) and seconds (horizontal axis).

calculates an estimate for the rate of change of the plasma's vertical velocity, weighted by the plasma's current.[13] Because the plasma current is approximately constant, this should, in principle, give a reasonable estimate for the plasma's vertical velocity and has been very successful in controlling the vertical stability of JET plasmas. An additional recent approximation made in this calculation is the neglect of a number of magnetic measurements from the vicinity of the divertor that, in principle, should be included.[13] The measurements were neglected after the inclusion of the conducting divertor structure because they were affected by currents flowing in the vessel and by noise from the amplifier.[13] It is not clear exactly why the vertical velocity estimate differs from that calculated using EFIT, but the different method of calculation and the use of a smaller subset of signals do allow this possibility. The approximations mentioned above, and the coincidence between plasma oscillations predicted by EFIT with the observation of ELMs, enhanced Be II radiation, and ion losses, suggests the velocity estimated by EFIT is more likely to be correct.

On average, the minima in the radial inner, outer, and upper gaps, and the maxima in the X-point's vertical displacement upwards, all occur approximately 3 ms before an ELM is triggered (see Figure 5). Given the 8 ms period of the oscillation, it appears that the ELM is triggered near what would otherwise be a maximum of the plasma's motion, near a minimum in the X-point's vertical position and near a maxima for the inner, outer, and upper gaps.

Next consider the plasma's edge velocity (Figure 10). Four milliseconds before the ELM, there is a minimum in the inner, outer, and upper plasma edge velocities (outwards away from the plasma's centre), and a maximum in the velocity of the plasma's X-point upwards. Again given the 8 ms period of the ELM, the ELM appears to be triggered at what would otherwise be the maxima of the plasma's velocity downwards and of its inwards radial narrowing.

Changes to the plasma's equilibrium have the potential to drive instabilities.[6–9,14] Modelling may help to understand the qualitative mechanisms by which the plasma is destabilised,[9,14] but we caution that calculations will be complicated by the transport processes that will occur during the plasma oscillations; that cannot yet be accurately modelled. We also caution that the changes in plasma equilibrium we describe are very small, just a few percent, and equilibrium

FIG. 8. The plasma current, offset so that they are zero at the ELM time ($t = 0$). Top to bottom: the radial position (RC), radial velocity (VRC), vertical position (ZC), and vertical velocity (VZC) of the current centre. Circles are from a typical pulse (83 794), the solid line is an average across the pulse set, and the dashed lines are the average $\pm$ the standard deviation. The distances (vertical axis) are measured in metres, and the time (horizontal axis) is measured in seconds.
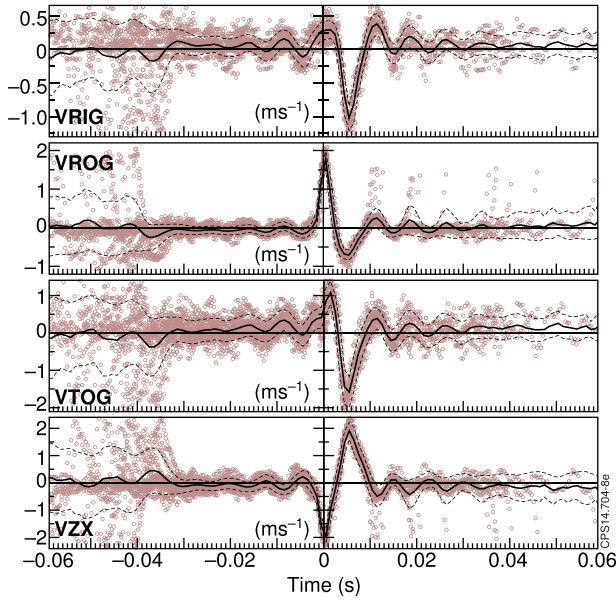
FIG. 10. Plasma edge velocity, offset so that they are zero at the ELM time ($t = 0$). From top to bottom: the velocity of plasma motion at the inner gap (VRIG), outer gap (VROG), plasma top (VTOG), and plasma's X-point (VZX). The absolute magnitude of the velocity depends on the numerical method by which the velocity is estimated, that requires the signal to be smoothed, but is in metres per second (vertical axis) with the horizontal (time) axis being in seconds. The velocities are slightly underestimated here, but the resulting signal shape is approximately independent of the smoothing procedure. Circles are from a typical pulse (83 794), the solid line is an average across the pulse set, and the dashed lines are the average ± the standard deviation.

reconstructions will in general be insufficiently accurate to allow conclusive inferences about changes to the plasma's stability.

III. FURTHER EVIDENCE FOR VERTICAL OSCILLATIONS

Probably the simplest and clearest evidence for a vertical oscillation of the plasma is from voltages induced in toroidal loops that circle above and below the plasma at an identical radius of 2.6 m from the central axis and are an approximately equal distance of ∼2.3 m above (or below) the plasma's centre. The voltage in each toroidal "flux loop" is plotted in Figure 11. A vertical plasma motion would be expected to produce equal but opposite electromotive forces (EMFs) in the top and bottom loops, with a slightly smaller EMF in the bottom loop due to the slightly greater shielding by conducting components near the divertor, and this is what is observed. It is difficult to imagine how a plasma perturbation could produce an EMF that is π out of phase and of similar period and magnitude unless it is due to a vertical oscillation, or how such an EMF could be produced by a control system without itself causing a motion of the plasma. The most likely explanation, that is consistent with all the other evidence here, is that the EMFs are due to a vertical plasma oscillation.

We can estimate the EMFs that will be induced in the toroidal flux loops due to the plasma's motion, and compare it with those observed. Consider an axisymmetric conducting loop with major radius R that encloses a circular area
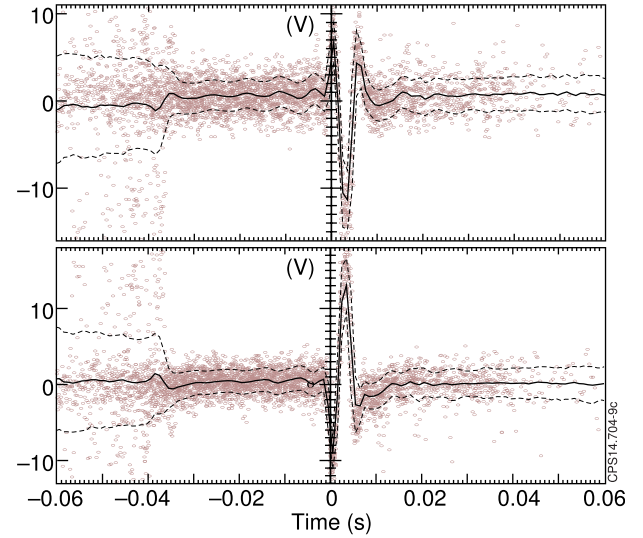

FIG. 11. The voltage (vertical axis, V) is plotted versus time (horizontal axis, s), for voltages measured by toroidal loops above the plasma (top figure), and below the plasma (bottom figure), with a positive voltage corresponding to the same direction as the plasma current and no-offset at ($t = 0$). Circles are from a typical pulse (83 794), the solid line is an average across every pulse in the set, and the dashed lines are the average ± the standard deviation. Oscillations of order 1 V in amplitude with an 8 ms period are very clear prior to ELMs (at $t = 0$) and are almost exactly π out of phase between the top and bottom plots, and of similar amplitude. This would be expected to be observed if the plasma was oscillating vertically with an 8 ms period.

$A = \pi R^2$. Using Stoke's theorem and Faraday's law ($\nabla \times \vec{E} = -\partial \vec{B} / \partial t$), the loop voltage per turn of conductor around the tokamak's axis is given by

$$\oint \vec{E} \cdot \vec{dl} = -\frac{\partial}{\partial t} \int_A \vec{B} \cdot \vec{dA}. \tag{1}$$

Taking the plasma current and the field it produces to be approximately independent of perturbations to the plasma's vertical position (although, in principle, there could be non-negligible modifications to the near-edge plasma and its stability). Then, the dependence of loop voltage on magnetic flux can be obtained from a Taylor expansion in terms of the perturbation δZ to the plasma's vertical position, giving the change in magnetic flux as

$$\delta \int_A \vec{B} \cdot \vec{dA} = \delta Z \frac{\partial}{\partial Z} \int_A \vec{B} \cdot \vec{dA}\Big|_{\delta Z=0} + O(\delta Z^2). \tag{2}$$

Dividing Eq. (2) by δt and taking $\delta t \to 0$, then gives

$$\frac{\partial}{\partial t} \int_A \vec{B} \cdot \vec{dA} = \left(\frac{dZ}{dt}\right) \frac{\partial}{\partial Z} \int_A \vec{B} \cdot \vec{dA}\Big|_{\delta Z=0} + \cdots. \tag{3}$$

A simple estimate for $\partial / \partial Z \int_A \vec{B} \cdot \vec{dA}$ can be found from dimensional analysis, with $\frac{\partial}{\partial Z} \int_A \vec{B} \cdot \vec{dA}\Big|_{\delta Z=0} \sim \frac{B_p}{Z_0} \pi R^2$ where B_p is of order the poloidal magnetic field in the flux loop, R is the major radius of the conducting loop, and Z_0 is the vertical distance to the current centre.[21]

For the plasmas considered in this paper, e.g., JET pulse 83 794, we take $B_p \sim 0.3$ T, $Z_0 \sim 2$ m, and Figure 10 shows maximum values for the pre-ELM oscillations of dZ/dt $\sim \pm 0.2$ ms^{-1}. With Eqs. (1) and (3), and the estimates

discussed above, we get, $\oint \vec{E} \cdot \vec{dl} \sim 0.1R^2$. The flux loops we consider are at $R = 2.6$ m, giving a maximum loop voltage of order ± 0.6 V/turn. The estimate is a similar order of magnitude to the ± 1 V pre-ELM oscillations shown in Figure 11.

A phase relationship has previously been observed between ELM occurrence times and the voltage measured by the flux loops VLD2 and VLD3 that circle toroidally around the tokamak either side of the divertor.[15] The cause of these now seems clear—it is the EMF induced by the plasma's oscillatory motion.

IV. SUMMARY AND CONCLUSIONS

120 almost identical JET H-mode plasmas have provided two orders of magnitude more data than is usually available from a single pulse, providing thousands of statistically equivalent ELMs. By combining this data, an unprecedentedly accurate description of the plasma's behaviour has become possible. The first surprising result was an observed sequence of maxima and minima in the probability density function for the waiting times between ELMs (Figure 1, reproduced from Ref. 3). Here, we also find a similar sequence of 8 ms-period maxima and minima in Be II light emissions, and in ion losses from the plasma's edge as determined from Langmuir probe and line-integrated density measurements. The maxima and minima's times coincide with the most (and least) likely waiting times for ELMs (the maxima and minima of the pdf in Figure 1). The evolution of the plasma's shape as inferred by EFIT[10,11] has been described in detail. Between ELMs the plasma moves diagonally radially outwards and downwards towards the divertor, with its position being rapidly reset back towards its starting point by the occurrence of an ELM. In addition to this comparatively slow diagonal drift of the plasma, EFIT predicts an oscillation that narrows the plasma radially and moves it vertically downwards as it narrows with a period of approximately 8 ms (125 Hz). Equal and opposite (π out of phase) EMFs that are measured by flux loops vertically above and below the plasma, seem to confirm the presence of a vertical plasma oscillation, and estimates for the EMFs that the plasma's motion would be expected to induce are consistent with those observed. These measurements provide the clearest evidence for plasma motion—we are unaware of any plasma instability that could produce the top and bottom loop voltages' phase, period, and amplitude, they can only be produced by a plasma motion or a vertical control system that itself would cause a plasma motion. Interestingly, the estimate for the plasma's velocity that is used by the vertical control system does not appear to show the vertical oscillations that are clearly shown in EFIT's calculations. This would explain why the oscillations are not being observed (and subsequently prevented) by the vertical control system. Because the oscillations found by EFIT coincide with enhanced (or reduced) probability of ELMs, with enhanced (or reduced) ion losses and Be II emissions and with flux loop measurements, it seems likely that EFIT's calculations are at least qualitatively correct. What is certain, is that EMFs with an 8 ms period are being observed by toroidal flux loops vertically above and below the plasma that are π

out of phase with each other, and that these coincide with enhanced (or reduced): Be II light emissions, ion losses from the plasma's edge, ELM occurrence, and EFIT predictions for the plasma's motion, with all of these using (or being verifiable by), independent diagnostics. Two important questions are as follows: (1) Is it easier (or more difficult) to trigger ELMs at frequencies that are in phase (out of phase) with the observed oscillation? (2) Can we deliberately trigger similar transport and instability processes by using the shape control system to produce similar oscillations to those found here to either (a) directly trigger ELMs or (b) enhance (or reduce) ion losses? If this was possible, it would be interesting to see whether the method would be able to enhance edge transport rates without triggering ELMs, with the ultimate aim of avoiding ELMs entirely, possibly in a similar way to "M-mode".[16] Both of these possibilities can be tested by experiment. Given the importance of these issues for plasma confinement in large tokamaks such as ITER, we hope that there will be an opportunity to explore these possibilities in the near future. Independent of this, the results suggest that small but possibly rapid changes in the plasma's position, and changes to edge-plasma transport and stability, cannot always be studied in isolation. For example, the oscillatory bursts of edge transport would not be captured by a transport model that did not identify vertical oscillations and correctly describe their consequences, and neither would a stability calculation correctly predict the onset of an ELM without also identifying that oscillations are occurring and correctly describing how they modify the plasma's stability. Undiagnosed changes to plasma density and position might partly explain the statistical nature of ELM occurrence that has been observed in both naturally occurring ELMs[3–5,17,18] and in experiments to deliberately trigger ELMs.[19] We hope that future work will explore all of these possibilities.

APPENDIX A: SYNCHRONISATION OF DATA

Plots, such as Figure 5, are of signals that are synchronised to the ELM times by plotting the signal between the $(n-1)$th and $(n+1)$th ELM, with the time t_n of the nth ELM subtracted from the time co-ordinate, and the value of the signal at the time of the ELM subtracted from the signal's value throughout. So time $t = 0$ corresponds to the

time of the nth ELM and has a zero value of the signal. Slicing up the signal in this way produces $(n-2)$ sets of signal data for a pulse with n ELMs, and the dataset contains roughly 120 pulses with roughly 40 ELMs in the 2.0 s between 11.5 s and 13.5 s. To display this data in a meaningful way, we average the data across many pulses in the dataset, plotting the averages from each of the pulses along with the average $\pm$ the standard deviation. This gives an indication of the average signal value and the scatter of the average between pulses. Because of the variation in waiting times between ELMs, the times $t_{n+1} - t_{n-1}$ vary for each triple of ELMs, and the time domain for which signal data will exist varies between each triple of ELMs over which it is sampled. Therefore, when averaging the signals, there is a choice of whether to average the signal by dividing by the number of these $(n-2)$ sets of signals or by the number of data values at each time point t to $t + dt$. The former is the average signal value observed before and after an ELM, the latter is the average signal value, given that there is data available at that value. The former is what on average you would expect to see, the latter tells you what you would expect the signal to look like if it is non-zero there. On this occasion, we have chosen to do the latter. We have also chosen to exclude pairs of ELMs with $t_n - t_{n-1}$ less than 40 ms, so as to try and exclude the strong post-ELM response that would be observed after the ELM at t_{n-1}, from the pre-ELM signal between $t_{n-1} - t_n$ and t$=0$ in the plots. Consequently, we are not studying any pairs of ELMs with $t_n - t_{n-1}$ less than 40 ms, this is unfortunate, but seems unavoidable if we wish to exclude the post-ELM response. We have considered longer and shorter periods, but 40 ms seems adequate to remove most evidence of the post-ELM response without throwing away too much data.

APPENDIX B: THE CENTRAL LIMIT THEOREM

Most of the plots, such as Figure 3, calculate averages across many individual ELMs. The JET plasmas were originally designed to be as similar as possible,[1] and we have only considered ELMs from the most stationary final 2 s of the H-mode phase of the plasma. Consequently, the ELMs are, in principle, statistically equivalent, and in practice, they are as close to being statistically equivalent as seems presently possible. The central limit theorem ensures that if you calculate an average or an average of a standard deviation across n events, then the error due to random noise falls like $1/\sqrt{n}$. There are of order 10 000 ELMs in these pulses,[3] although we often only need to consider a subset of them. For illustration, if we had a 10 cm random error in a measurement that was averaged over 10 000 ELMS, then the error would fall to $0.1/\sqrt{10000}$ m = 1 mm. Consequently, by averaging over many equivalent ELMs we are effectively eliminating random errors from measurements and calculations by codes. Systematic errors can still be present and can lead to different absolute values such as size of displacements, but the effective removal of random errors allows an accurate resolution of oscillations that would otherwise be within the error bars of a single measurement. In other words, by averaging over many thousands of statistically

equivalent events the random errors will become negligible, allowing detail to be resolved in the averages that can be smaller than the (random) error or variation in an individual measurement.

APPENDIX C: SELECTION OF DIAGNOSTICS AND SOURCES OF UNCERTAINTY

We have used diagnostics that provide independent information on the plasma's behaviour. The diagnostics that we could use were in practice limited by two factors in particular:

1. The rate at which the data were sampled—If too low it would poorly resolve the ELM event and pre-ELM oscillations, so attention was restricted to data that was frequently sampled, typically with a rate of one point every 0.1 or 0.2 ms.
2. The data needed to be sampled in an automatic way that allows large datasets to be processed computationally—this is essential to ensure the statistical increase in accuracy that can be accomplished by averaging over many pulses (if the data associated with roughly 10 000 ELMs are combined, then the random errors are, in principle, reduced by a factor of order 1/100).

If one or more people are needed to spend dedicated time (hours typically) to process a particular signal, then the data cannot be analysed on such a large scale, and the statistical gains in accuracy are lost. The larger error bars would then prevent a discussion of the small oscillations that we consider. Three further factors were as follows:

3. Whether the diagnostic was available; reflectometry, for example, does not appear to have been available for these pulses.
4. Whether the diagnostics are independent—there is no advantage in looking at a diagnostic output that relies on EFIT, for example, and then comparing it with EFIT's results.
5. There are periodic losses of ions from the plasma's edge (see Figure 4), which makes it very difficult to determine if changes in a line-integrated measurement through the plasma's top or bottom are due to changes in density or to a plasma motion that changes the line-integral's path length.

For the work presented here, and bearing in mind the remarks above, there was a thorough search for diagnostics that could provide reliable and independent diagnostic information on the plasma's behaviour. Results from those that were most useful according to the criteria discussed above are presented here.

Sources of uncertainty for EFIT: For the ELMs studied, the plasma is in a "flat top," during which any Halo currents due to (comparatively small) plasma motions are negligible compared with the control coil currents. Currents can also be induced in JET's toroidally conducting divertor structure, whose resistance was of order 3×10^{-3} Ω prior to the installation of the ITER-like wall and the replacement of carbon tiles with metal plasma-facing components. Its present

resistance is unfortunately unknown due to a failure of the diagnostic previously used to measure the currents that flow. However, taking the resistance of order 3×10^{-3} that was inferred from older measurements of induced divertor currents and loop voltages, the currents driven by the vertical oscillation's of order 1 V induced voltages (see Section III), would be a few hundred amperes. This is about 1% of the total combined current flowing in the nearby divertor coils, and about 0.01% of the plasma's current. So the currents could be influencing EFIT, and we would prefer for future measurements that are able to check this more thoroughly, but any influence is expected to be small and we are unable to test their influence at present. More importantly, both physical sources of error just described will be small compared to the inherent uncertainties in the magnetic reconstruction that can exaggerate small measurement uncertainties and typically range from a few percent to ten percent or more for some quantities.[20]

The clearest evidence for vertical plasma oscillations is the π out of phase toroidal loop voltages measured above and below the plasma, and the consistency between these, and all the other independent diagnostics. The consistency between these and EFIT's calculations is what provides the greatest source confidence that EFIT's results are approximately correct.

[1]S. Brezinsek, T. Loarer, V. Phillips *et al.*, Nucl. Fusion **53**, 083023 (2013).

[2]R. Neu, G. Arnoux, M. Beurskens *et al.*, Phys. Plasmas **20**, 056111 (2013).

[3]A. J. Webster, R. O. Dendy, F. A. Calderon, S. C. Chapman, E. Delabie, D. Dodt, R. Felton, T. N. Todd, V. Riccardo, B. Alper, S. Brezinsek, P. Coad, J. Likonen, M. Rubel, and JET EFDA Contributors, Plasma Phys. Controlled Fusion **56**, 075017 (2014).

[4]A. J. Webster and R. O. Dendy, Phys. Rev. Lett. **110**, 155004 (2013).

[5]A. J. Webster, S. J. Webster, and JET EFDA Contributors, Phys. Plasmas **21**, 112502 (2014).

[6]A. W. Degeling, Y. R. Martin, J. B. Lister *et al.*, Plasma Phys. Controlled Fusion **45**, 1637 (2003).

[7]P. T. Lang, A. W. Degeling, J. B. Lister, Y. R. Martin, P. J. Mc Carthy, A. C. C. Sips, W. Suttrop, G. D. Conway, L. Fattorini, O. Gruber, L. D. Horton, A. Herrmann, M. E. Manso, M. Maraschek, V. Mertens, A. Mück, W. Schneider, C. Sihler, W. Treutterer, H. Zohm, and ASDEX Upgrade Team, Plasma Phys. Controlled Fusion **46**, L31 (2004).

[8]F. Sartori, P. Lomas, F. Piccolo, M. K. Zedda, and JET EFDA Contributors, in *35th EPS Conference on Plasma Physics, Hersonissos, 9–13th June* (ECA, 2008), Vol. 32D, p. P-5.045.

[9]Y. Liang, Fusion Sci. Technol. **59**, 586 (2011).

[10]L. L. Lao, H. St. John, R. D. Stambaugh, A. G. Kellman, and W. Pfeiffer, Nucl. Fusion **25**, 1611 (1985).

[11]L. C. Appel, G. T. A. Huysmans, L. L. Lao *et al.*, in *33rd EPS Conference on Plasma Phys. Rome, 19–23 June 2006* (ECA, 2006), Vol. 30I, p. P–2.184.

[12]F. Sartori, G. De Tommasi, and F. Piccolo, IEEE Control Syst. Mag. **26**, 64–78 (2006).

[13]R. Albanese, T. Bellizio, and JET-EFDA Contributors, Fusion Sci. Technol. **59**, 363 (2011).

[14]A. J. Webster, Nucl. Fusion **52**, 114023 (2012).

[15]S. C. Chapman, R. O. Dendy, T. N. Todd, N. W. Watkins, A. J. Webster, F. A. Calderon, J. Morris, and JET EFDA Contributors, Phys. Plasmas **21**, 062302 (2014).

[16]E. R. Solano, N. Vianello, P. Buratti *et al.*, in *40th EPS Conference on Plasma Physics* (2013), p. P4.111.

[17]A. Murari, F. Pisano, J. Vega *et al.*, Plasma Phys. Controlled Fusion **56**, 114007 (2014).

[18]A. Shabbir, G. Verdoolaege, OJWF Kardaun *et al.*, Rev. Sci. Instrum. **85**, 11E819 (2014).

[19]A. J. Webster and JET EFDA Contributors, Phys. Plasmas **21**, 112506 (2014).

[20]L. L. Lao, J. R. Ferron, R. J. Groebner, W. Howl, H. St. John, E. J. Strait, and T. S. Taylor, Nucl. Fusion **30**, 1035 (1990).

[21]A few notes about this estimate: (i) A calculation using a circular cross-section plasma confirms this. (ii) In principle, $\partial/\partial Z \int_A \vec{B} \cdot \vec{dA}|_{\delta Z=0}$ can be calculated exactly, but the estimates here are sufficiently accurate for the purposes of this paper. (iii) In principle, $\partial/\partial t \int_A \vec{B} \cdot \vec{dA}|_{\delta Z=0}$ can be calculated or estimated from experimental measurements. (iv) A caveat—for radii R that are larger than the radial position of the plasma's current centre, contributions can be in equal and opposite directions and will partially cancel. For those cases, the estimates will be less reliable.