INSTITUTE OF PHYSICS PUBLISHING PLASMA PHYSICS AND CONTROLLED FUSION

Plasma Phys. Control. Fusion **45** (2003) 747–758 PII: S0741-3335(03)54112-4

# Probability distribution functions for ELM bursts in a series of JET tokamak discharges

**J Greenhough[1], S C Chapman[1], R O Dendy[1,2] and D J Ward[2]**

[1] Space and Astrophysics Group, Department of Physics, Warwick University, Coventry CV4 7AL, UK
[2] EURATOM/UKAEA Fusion Association, Culham Science Centre, Abingdon, Oxfordshire OX14 3DB, UK

E-mail: greenh@astro.warwick.ac.uk and richard.dendy@ukaea.org.uk

Received 1 October 2002, in final form 10 February 2003
Published 4 April 2003
Online at stacks.iop.org/PPCF/45/747

**Abstract**

A novel statistical treatment of the full raw edge localized mode (ELM) signal from a series of previously studied JET plasmas is tested. The approach involves constructing probability distribution functions (PDFs) for ELM amplitudes and time separations, and quantifying the fit between the measured PDFs and model distributions (Gaussian, inverse exponential) and Poisson processes. Uncertainties inherent in the discreteness of the raw signal require the application of statistically rigorous techniques to distinguish ELM data points from background, and to extrapolate peak amplitudes. The accuracy of PDF construction is further constrained by the relatively small number of ELM bursts (several hundred) in each sample. In consequence the statistical technique is found to be difficult to apply to low frequency (typically Type I) ELMs, so the focus is narrowed to four JET plasmas with high frequency (typically Type III) ELMs. The results suggest that there may be several fundamentally different kinds of Type III ELMing process at work. It is concluded that this novel statistical treatment can be made to work, may have wider applications to ELM data, and has immediate practical value as an additional quantitative discriminant between classes of ELMing behaviour.

## 1. Introduction

The central role of edge localized modes (ELMs) in the physics of toroidal magnetic fusion plasmas is well known. ELM physics and theory are reviewed, for example, in [1], and the practical importance of ELMs for present-day tokamak operation and power plant design are further emphasized in brief reviews oriented towards power and particle exhaust [2] and towards confinement regimes [3]. ELMs are also found in association with improved confinement regimes in the W7-AS stellarator [4]. Key ELM issues for the ITER H-mode, for example, include both wall loading and the effect on coupling of ion cyclotron heating antennas to the

  

plasma. ELM phenomenology in the joint European torus (JET), with which this paper is primarily concerned, is well documented (see, e.g. [5, 7]).

Theoretical interpretation of ELMs has been attempted primarily through magnetohydrodynamic (MHD) models. For example, ELM excitation may be caused [8] by the interplay between steep edge pressure gradients, local magnetic shear, and the bootstrap and Pfirsch–Schlüte currents, leading to linear instability of pressure-driven ballooning modes and current-driven peeling modes. There are experimental indications that in some cases the character of ELMs may depend, *inter alia*, on proximity to MHD stability boundaries; for a discussion, see [1] and references therein. While it is important to identify the linear MHD instabilities that may underlie the initial stage of ELM excitation, the nonlinear and bursty character of ELMs is difficult to capture using an MHD approach. This motivates two new channels of research. First, one may ask whether ELM phenomenology is unique to magnetic fusion plasma systems. The answer appears to be 'no'. It has recently been shown [9] that the scaling of frequency of avalanche pulses involving mass loss and rapid inward edge movement from a particular sandpile model [10] (one which self-organizes steep edge gradients) with stored energy is in a way analogous to the scaling of ELM frequency with stored energy in some types of JET plasma. This outcome lends support to a second incipient channel of research, which we pursue in this paper: applying to ELM data, some of the other recent developments in statistical physics and time series analysis (see, e.g. [11]). Previous work [12, 13] on applications to ELMs focused on investigating whether unstable periodic orbits can be identified in the time series of intervals between ELM pulses in certain plasmas in the JT-60U and TCV tokamaks. However, we find no correlation between burst amplitudes and intervals, for the ELM pulses from the series of JET plasmas considered here, and techniques such as delay plots do not reveal a low dimensional phase space. There is, of course, no expectation that all ELMs in all tokamaks should have the same properties when evaluated as nonlinear time series. Here, therefore, we adopt a different and conceptually simple statistical approach, involving the construction and analysis of probability distribution functions (PDFs) for ELM bursts for our series of JET discharges.

It is worth recording that it is always challenging to draw statistically meaningful inferences from tokamak data, especially in comparison to the cleaner datasets obtainable from smaller scale experiments in other branches of physics. For example, Wegend and Gershenfeld [14] describe a variety of physical (and other) time series which are well adapted to the application of modern techniques, in part, because they are quasistationary and include very large numbers of events. In contrast, the difficulty in identifying substantial episodes of quasistationarity yielding long clean datasets during tokamak discharges is well known. Conversely, tokamak plasmas provide an interesting challenge for the practical application of modern time series analysis techniques to complex large scale physical systems. This paper is thus primarily concerned with proof-of-principle: we want to provide an indication of how useful modern techniques may be, in the tokamak context. For some of the JET plasmas considered here, it has indeed proven possible to construct PDFs for the ELMs that enable clear conclusions to be drawn. This provides a new quantitative method for distinguishing ELM characteristics. A key practical result is that the PDFs of time intervals between ELMs indicate that ELM processes with fundamentally different statistical–physics properties can operate in different discharges.

## 2. Characterization of the full ELM signal

In this paper, we take a fresh look at ELMs, from a perspective which may be unfamiliar. Much has been learnt about ELMs by classifying them on phenomenological grounds into

Type I, Type III, and other groups. In particular, the ELMing JET plasmas considered in this paper have previously been studied in depth in [15], to which we refer for experimental detail, plasma characterization, and a description of the ELM phenomenology in conventional terms. Here, however, we return to the raw ELM time series obtained from these plasmas, and seek to evaluate their statistical properties directly, thereby testing the applicability of modern approaches [11]. From the wide range of JET plasmas considered in [15], we have identified four that have quasistationary episodes during which there occur enough ELMs for statistically meaningful PDFs to be constructed. Figure 1 shows ELM traces from the quasistationary phases of these four JET plasmas, which differ primarily in their isotopic concentration [15]. Each trace contains approximately 50 000 data points, sampled at 0.1 ms intervals, and includes several hundred ELM bursts.

Figure 1(*a*) shows the complete ELM trace from beginning to end of JET plasma 43002, with the approximately quasistationary phase (identified by eye) outlined by dashed lines. The trace during this phase is expanded in figure 1(*b*). Figures 1(*c*)–(*e*) were obtained in the same way as figure 1(*b*), for JET plasmas 43392, 44028, and 44349. It is the data imaged in figures 1(*b*)–(*e*) that is analysed further in this paper. We note that while the total number of data points is large, the number of ELM bursts is not, and this represents one of the most challenging aspects of the data. With only (say) 500 ELM bursts in a given discharge, it may be possible to confirm or exclude low-dimensional behaviour (for example quasiperiodicity), but it will be difficult to confirm or characterize high dimensional behaviour (which fills much of phase space) unambiguously. Our objective here is therefore to focus on these traces as time series. While it is clear from inspection that the trace for JET plasma 43392 (figure 1(*c*)) stands out, it is interesting to see how far one can place this observation on a quantitative statistical basis. Non-stationarity of the data imposes an additional constraint. In particular, it is clear from figure 1(*b*) that there is a transition away from the steady (Type III) behaviour after approximately 1 s in the ELM signal from JET plasma 43002. While it would be desirable to form PDFs for this steady state alone, such analysis is not possible because there are only 28 bursts in this section. Likewise, the PDFs that we plot below (section 3) for the whole of series 43002 (comprising 118 bursts) are not significantly altered by the removal of the initial 1s. We are therefore unable to rigorously quantify the apparent transition in behaviour in this emission, and use the whole series in the analyses that follow.

The data points in these time series contain information about ELM bursts and about activity during the intervals between bursts—if it is justifiable to assume that these two classes of contribution to the signal are well defined and separable. In this section we address two necessary preliminaries, so that the subsequent statistical analysis of these ELM burst distributions can be developed on a firm basis. First, we need to quantify how clearly ELM bursts can be distinguished from interburst activity in the raw time series. Second, we need to construct an algorithm for isolating and quantifying individual ELM bursts. This algorithm will also need to deal with the discrete nature of the sampling, which means that the actual maximum point of each ELM burst is likely to be missing from the time series.

Figure 2 plots the power spectra of the full ELM signal time series that are displayed in figure 1. Unlike figures 3 and 4 in the next section where, as described below, the data has been thresholded so as to define isolated ELM bursts, figure 2 contains information about both the ELM bursts and the higher frequency signal within which the bursts reside. Let us recall that the Fourier transform of an $N$-point discrete time series $x_n$ is given by

$$p_k = \frac{1}{\sqrt{N}} \sum_{n=1}^{N} x_n \, \mathrm{e}^{2\pi \mathrm{i} k n/N}$$

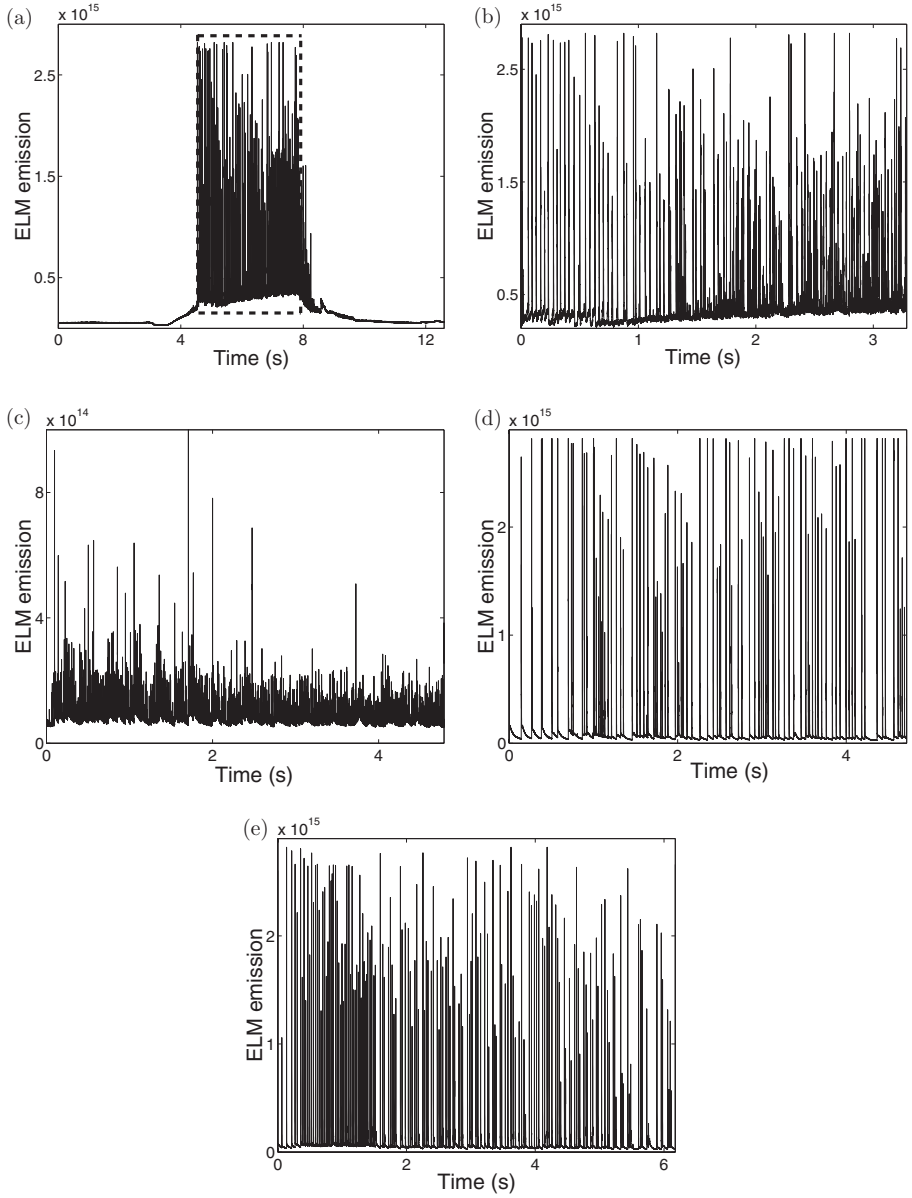


**Figure 1.** (*a*) Time series of ELM signal throughout the duration of JET plasma 43002, with the visually identified quasistationary phase shown. (*b*)–(*e*) Time series of ELM emission (arbitrary units) during quasistationary phases of JET plasmas 43002, 43392, 44028 and 44349, which differ primarily in isotopic concentration. Sampling interval 0.1 ms.

where the frequencies in physical units are $f_k = f_{\text{samp}}k/N$ for $k = 0, \ldots, N/2$ and $f_{\text{samp}}$ is the sampling frequency. The power spectrum is defined to be the square of the amplitudes with which frequencies $f_k$ contribute to the signal: $P_k = |p_k|^2$. To assist interpretation, superimposed on each spectrum in figure 2 is a vertical dashed line located at the inverse of the mean inter-burst time interval measured between the ELM bursts in that discharge, which we denote $\nu_{\text{ELM}}$. Typical values for $\nu_{\text{ELM}}$ are in the range 30–100 Hz. The clear (but broad) peaks

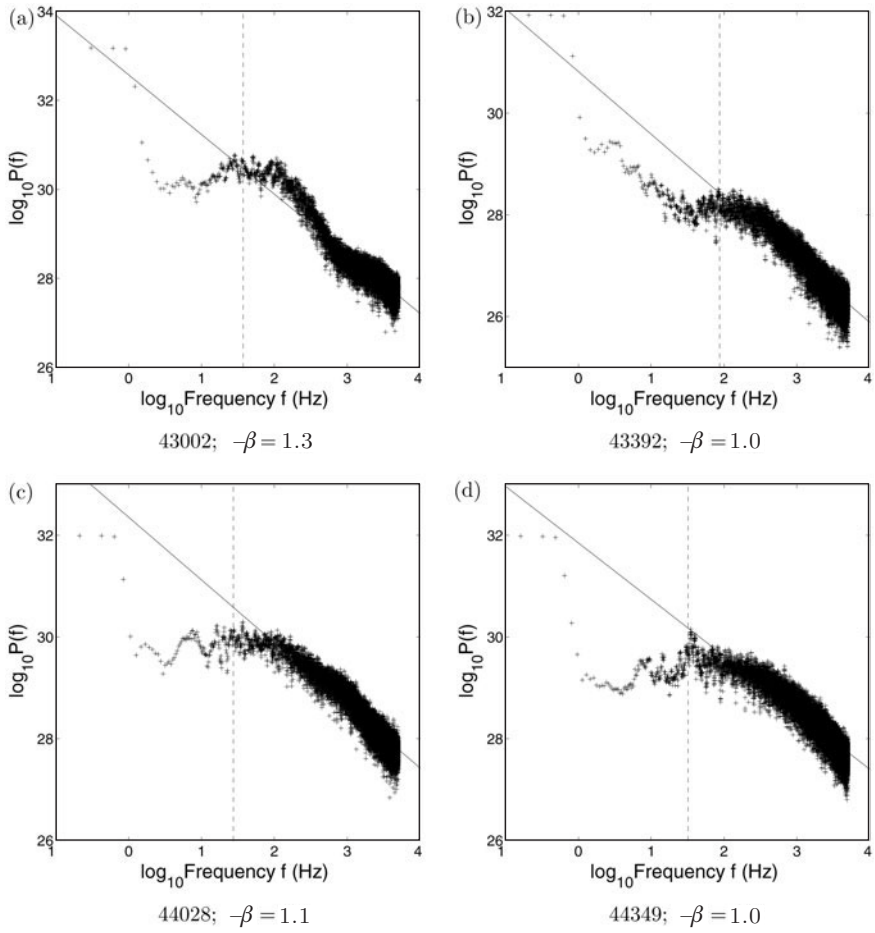


**Figure 2.** Logarithmic plots of ELM emission power spectra. Least-squares lines of best fit have slope $-\beta$ as given in the sub-captions; dashed lines indicate inverse mean inter-burst interval.

visible below $\nu_{\mathrm{ELM}}$ in the power spectra of figures 2($c$) and ($d$) correspond to the episodes of regularly spaced ELMing that can be seen in figures 1($d$) and ($e$).

One can define (see, e.g. [16]) the degree of persistence $\beta$ in the signal as minus the slope of the best fit straight line on a plot of $\log P_k$ against $\log k$, or equivalently a plot of $\log P(f)$ against $\log f$ (for an example of an application see [17]). We apply Thompson multi-tapering to reduce the variance [18] before fitting a least-squares straight line of slope $-\beta$. The values of $\beta$ thus obtained have been confirmed by wavelet analysis.

Physically, the power spectra in figure 2 imply that, to first approximation, the high frequency activity that persists between ELM bursts shows persistence ($\beta \approx 1$ for $\nu \geqslant \nu_{\mathrm{ELM}}$). Here again we are operating at the statistical limits imposed by the nature of the data and by the physics of the ELM bursts. Since the power spectrum is averaged over all times, and the inter-burst intervals are not equal, there is no abrupt transition from the persistent noise between bursts, to the broadly quasiperiodic behaviour on longer timescales. It remains clear, however, that there is a qualitative transition, localized in the frequency domain near $\nu_{\mathrm{ELM}}$.

These results quantify the extent to which the ELM bursts constitute a statistically distinct component of the raw signal in the JET discharges considered here. They also confirm that

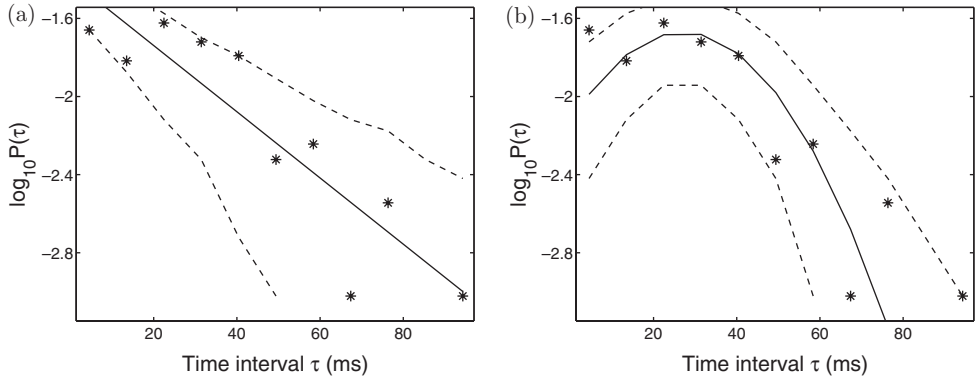


**Figure 3.** PDF of time intervals $\tau$ between ELM bursts for JET plasma 43002, showing: (*a*) best fit inverse exponential PDF fitted by least squares; (*b*) a Gaussian model with the same $\mu$ and $\sigma$ as the data. Dashed lines indicate upper and lower confidence limits inferred from 99% binomial counting errors. Computed values for the goodness-of-fit parameter $R^2$ are (*a*) 0.81 and (*b*) 0.58.

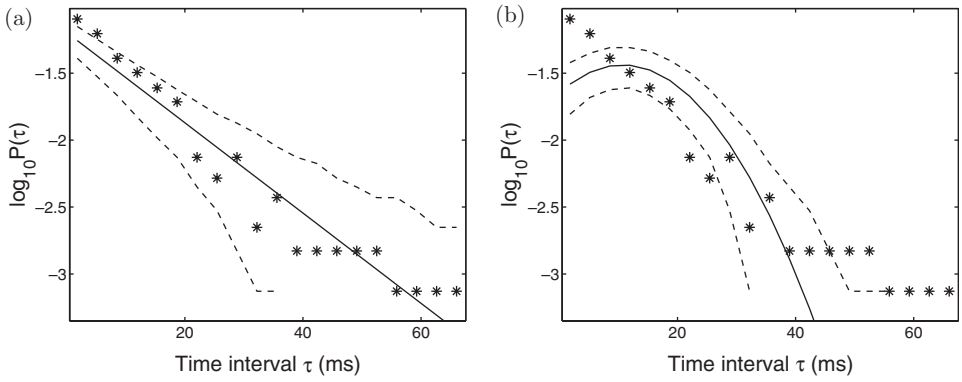


**Figure 4.** PDF of time intervals $\tau$ between ELM bursts for JET plasma 43392, showing: (*a*) best fit inverse exponential PDF fitted by least squares; (*b*) a Gaussian model with the same $\mu$ and $\sigma$ as the data. Dashed lines indicate upper and lower confidence limits inferred from 99% binomial counting errors. Computed values for the goodness-of-fit parameter $R^2$ are (*a*) 0.93 and (*b*) 0.31.

it is legitimate to construct a thresholding process to define and quantify the individual ELM bursts. We therefore turn to this second necessary preliminary to computing and analysing the statistical properties of ELM bursts through their PDFs.

As noted above, the discrete nature of the sampling means that the exact magnitude of the maximum point of each ELM burst is likely to be missed. To estimate the true maximum amplitude from the discrete time series, we assume an exponential decay from the maximum, and take a least-squares log-linear fit to the observed maximum and the subsequent point. This is extrapolated back one timestep from the observed maximum, to give the inferred maximum amplitude. The beginning of an ELM burst is defined as the point whose magnitude is smaller than that at the succeeding point by a factor $J > 1$, whose value is chosen empirically. We find that $J = 2$ selects those events that one would identify as ELM bursts by visual inspection. Thus, our thresholded ELM characterization algorithm models the true ELM burst from the data, in terms of an instantaneous rise from its starting point to the estimated maximum, followed by exponential decay over two time steps. The PDFs for ELM burst properties in

a given discharge are then obtained by binning the resulting data into $n^{1/2}$ bins, where $n$ is the number of ELM bursts identified. This provides a compromise between resolution and counting fluctuations. Normalization of the PDFs to unity is achieved by dividing the number of points in each bin by $n$ times the binwidth.

## 3. Statistical distributions of ELM bursts

For the four JET plasmas 43002, 43392, 44028 and 44349 whose ELM signals are shown in figure 1, it is possible to construct PDFs of the measured distributions of ELM amplitudes, and of the measured distributions of time intervals $\tau$ between ELMs. These PDFs are useful insofar as their properties shed light on the underlying character of the ELMing process in the different plasmas. Here, we focus on testing two leading hypotheses:

- The PDFs are Gaussian, implying that the measured ELM properties are randomly distributed about characteristic (mean) values.
- The PDFs are inverse exponential. In the time domain, for example, such a distribution would arise from a Poisson process where events are independent and the event probability per unit time is constant.

As we show below, for our measured ELM PDFs it is not always possible to distinguish between these hypotheses with a high level of confidence. Nevertheless, in some cases firm conclusions can be drawn, and in all cases, the study of the PDFs provides a new quantitative perspective on ELM physics. Our tests of these hypotheses begin by constructing best fit curves: Gaussians are plotted using the measured mean $\mu$ and standard deviation $\sigma$ of the PDF; exponentials are fitted by least squares. We evaluate the goodness-of-fit between the measured PDFs and the fitted Gaussian or inverse exponential model curves in two ways, as follows.

Graphically, we plot curves on both sides of the best fit which show upper and lower 99% binomial counting errors. In figures 3–6, for example, the dashed lines reflect estimates of the likely counting errors in the number of points in the bins, arising from the finite size of the dataset and assuming a binomial distribution of counting errors. We have computed the probability distribution of the number of points in each bin, given both the number $n$ of ELMs identified, and the expected value inferred from the fitted PDF: the dashed lines in figures 3–6 thus represent the limits between which the observed data are expected to fall on 99% of occasions.

Mathematically, we compute the square of the correlation coefficient $R$, where $1 - R^2$ is defined to equal the ratio of the sum of squared residuals from the model to the sum of squared deviations of the data from the mean [19]. $R^2$ can be described as the proportion of variation explained by the model (as opposed to explained by data variance), or the relative predictive power of a model over a constant (the mean), and hence is often referred to as the coefficient of determination. A limitation of $R^2$ is that it weights all deviations equally, regardless of their likelihood, so that it does not capture how likely the data are to be observed, given the model. In this respect the 99% confidence bands are more illuminating.

### *3.1. Time intervals between ELMs*

Figure 3 displays the PDFs of the time intervals $\tau$ between ELMs for JET plasma 43002, with best fit inverse exponential and Gaussian curves plotted in figures 3(*a*) and (*b*), respectively. It is unclear which fitted curve provides a better model for the data. On one hand, the number of points that lie close to the Gaussian curve, especially near its peak, is substantially greater

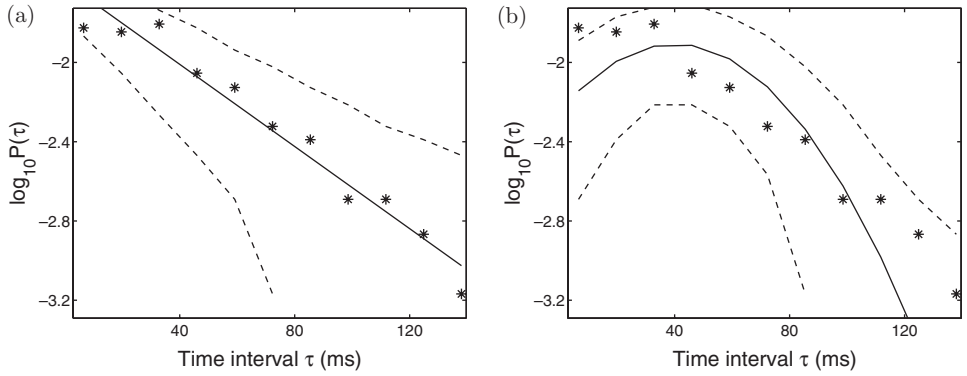


**Figure 5.** PDF of time intervals $\tau$ between ELM bursts for JET plasma 44028, showing: (*a*) best fit inverse exponential PDF fitted by least squares; (*b*) a Gaussian model with the same $\mu$ and $\sigma$ as the data. Dashed lines indicate upper and lower confidence limits inferred from 99% binomial counting errors. Computed values for the goodness-of-fit parameter $R^2$ are (*a*) 0.96 and (*b*) 0.54.

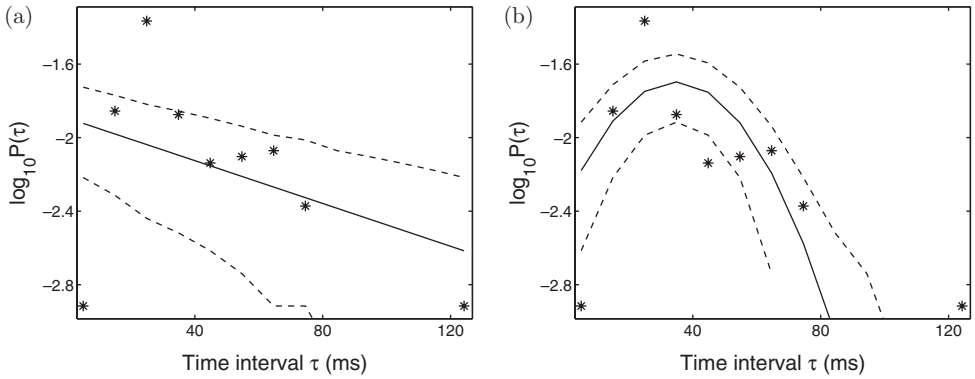


**Figure 6.** PDF of time intervals $\tau$ between ELM bursts for JET plasma 44349, showing: (*a*) best fit inverse exponential PDF fitted by least squares; (*b*) a Gaussian model with the same $\mu$ and $\sigma$ as the data. Dashed lines indicate upper and lower confidence limits inferred from 99% binomial counting errors. Computed values for the goodness-of-fit parameter $R^2$ are (*a*) 0.18 and (*b*) 0.26.

than the number that lie close to the exponential curve. On the other hand, the computed value of $R^2$ is 0.81 for the exponential, but only 0.58 for the Gaussian.

For JET plasma 43392, there are strong indications that the inverse exponential curve (figure 4(*a*)) is a better fit than the Gaussian (figure 4(*b*)). The graphical results of figure 4 are supported by the computed value of $R^2$, which is 0.93 for the exponential, but only 0.31 for the Gaussian.

Similarly, for JET plasma 44028, there are also strong indications that the inverse exponential curve (figure 5(*a*)) is a better fit than the Gaussian (figure 5(*b*)). The computed value of $R^2$ is 0.96 for the exponential, but only 0.54 for the Gaussian.

Figure 6 indicates that for JET plasma 44349, neither curve is a good fit. This is confirmed by low computed $R^2$ values: 0.18 for exponential and 0.26 for Gaussian.

It follows that the statistical properties of the distribution of time intervals $\tau$ between ELMs are more varied than might be inferred from visual inspection of the ELM traces of figure 1. This implies that fundamentally different statistical processes may be at work. So

far as we know, figures 3–6 are the first statistical characterization of ELM burst PDFs during a discharge, and yield the first statistical differentiation between ELM burst PDFs in different discharges. In particular, we have found that there is a good quantitative basis for inferring that the PDFs of ELM time intervals for JET plasmas 43392 and 44028 follow an inverse exponential distribution. This result motivates the investigation in section 3.2 of whether these reflect an underlying Poisson process.

### *3.2. Test for Poisson process*

In a Poisson process, there is a constant probability per unit time of an event occuring. Inverse exponential PDFs for the intervals between events arise naturally from Poisson processes, which suggests that the ELMs in JET plasmas 43392 and 44028 may arise from a Poisson process. To test this hypothesis, we recall the further implication that the number of bursts occuring in an arbitrary fixed time interval should be Poisson distributed. In figure 7, therefore, we plot the PDFs for the number $n$ of ELM bursts occuring within a fixed time interval $T$, in (for completeness) all four of the JET plasmas. The value of $T$ for each plasma is chosen to

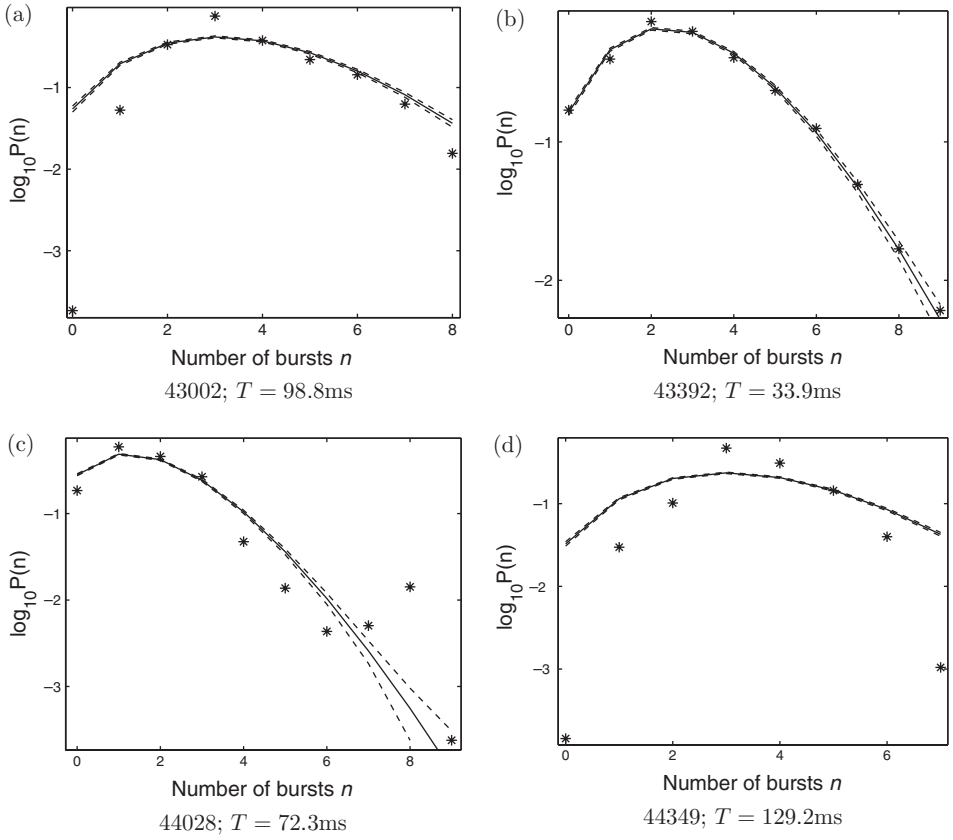


**Figure 7.** PDFs for the number $n$ of ELM bursts occuring within a fixed time interval $T$. Solid lines are Poisson PDFs that have the same mean $\mu$ as the data; dashed lines are for 99% binomial counting errors. Only the numbers of bursts in JET plasma 43392 are consistent with a Poisson distribution. Computed values for the goodness-of-fit parameter $R^2$ are 0.23 for 43002, 0.72 for 43392, 0.15 for 44028 and 0.06 for 44349.

be sufficiently large for a range of values of $n$ to be observed, but not so large that a Poisson distribution will approach its Gaussian high-mean limit (a Gaussian PDF would be parabolic on these plots). The solid lines are Poisson PDFs that have the same mean $\mu$ as the data, and the closely adjacent dashed lines are for 99% binomial counting errors. It is clear from figure 7 that only the distribution of bursts in JET plasma 43392 (figure 7(*b*)) is consistent with a Poisson process: its value of $R^2$ is 0.72, compared to the very small values (0.23, 0.15, and 0.06) for the other three JET plasmas.

This test thus also provides an interesting further level of quantitative discrimination between the ELMing processes in the JET plasmas considered here. Despite the fact that the ELM interval PDFs for JET plasmas 43392 and 44028 are fitted by inverse exponential distributions with similar high values of $R^2$ (0.93 and 0.96, respectively), their Poisson fits are radically different (0.72 compared to 0.15, respectively).

### *3.3. Distribution of ELM amplitudes*

The PDFs of ELM burst amplitudes are shown in figure 8. In contrast to the inter-burst time interval PDFs of figures 3–6, the degree of confidence $R^2$ with which either Gaussian or inverse exponential curves can be fitted to the amplitude PDFs is relatively low. Clear differentiation can be made only for JET plasma 44349. Here $R^2$ takes the value 0.77 for the best Gaussian fit, but only 0.04 for the best inverse exponential. Thus the PDFs of ELM $\tau$ enable statistical

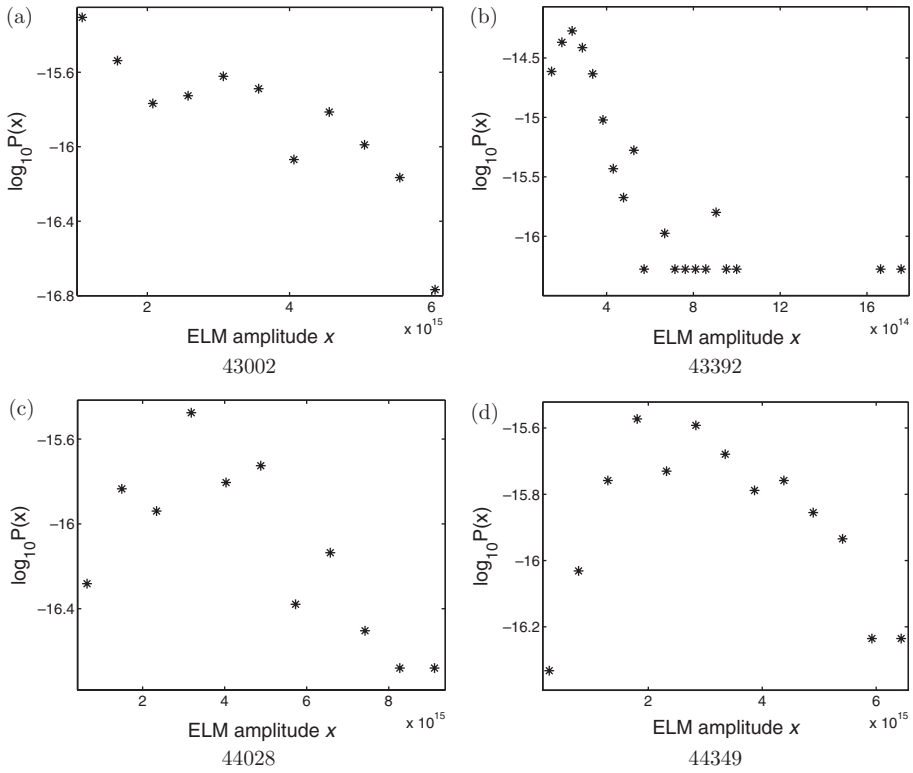


**Figure 8.** PDFs of ELM burst amplitudes (arbitrary units). For 44349, a Gaussian curve with the same $\mu$ and $\sigma$ as the data fits the PDF with $R^2 = 0.77$. This is the only convincing fitting however, and for this reason fitted curves are not shown.

differentiation between the JET plasmas considered, to a greater extent than do the PDFs of ELM amplitude.

We also note that an inverse exponential distribution of inter-burst intervals $\tau$ is compatible with the original sandpile paradigm [20] for self-organized criticality (SOC). However, a power-law PDF of burst amplitudes is also a necessary condition for SOC [11], and figure 8 shows that this is not satisfied for the series of JET plasmas considered here.

## 4. Discussion: implications for ELM physics

Thus far, we have treated the ELM data purely as time series whose statistical characteristics are to be established, without referring to the plasmas in which the observations were made. In fact the plasmas analysed are those produced in a systematic study [15] of ELM behaviour in JET under varying plasma conditions, in which changes to the plasma shape, isotopic composition and heating power were used to vary the phenomenological character of the ELMs. Although we have been concerned primarily with the applicability and overall results of the analysis, rather than a survey of ELM types, let us now link the results back to the corresponding plasmas, particularly with regard to the standard ELM categorization. In this study we have shown that the ELM data considered fall into two groups, distinguished by the nature of the PDF of waiting time $\tau$ between ELMs. In addition JET plasma 43392 is singled out by our analysis as being compatible with an underlying Poisson process. One might expect this grouping to correspond to that often used in fusion plasma physics, with a discrimination between Type I and Type III ELMs. However, what we find does not correspond to this. In fact the time series from JET plasmas 43392 and 43002 are both identified in [15] as Type III ELMs, but are nonetheless distinguished by the present analysis as having different statistical characters. A possible explanation is that JET plasma 43392 (which corresponds to the hydrogen plasma shown in figure 9 of [15]) is only marginally above the H-mode threshold, whilst JET plasma 43002 is a more secure H-mode. This suggests a hypothesis that the present analysis might help distinguish behaviour close to and far from the H-mode threshold. This is a tentative suggestion based on the evidence of only a few JET plasmas, but clearly invites further attention.

In this paper we have taken several steps that should encourage and assist future statistical studies of the distributions of ELM bursts in different plasma discharges and in different tokamaks. We find that the correlation properties of the entire ELM signal imply that a rigorous statistical distinction between ELM bursts and inter-burst activity can be made for the JET plasmas [15] considered. The uncertainties inherent in the discreteness of the raw ELM signal time series are overcome using a thresholding algorithm, which in turn enables us to construct PDFs for the magnitudes and intervals of ELM bursts. In plasmas with low frequency (typically Type I) ELMs, the number of ELMs during quasistationary episodes is too small for useful statistical inferences to be drawn from the PDFs. However in the four JET plasmas considered that include substantial episodes of Type III ELMing, it is possible to make quantitative comparisons between fitted statistical functions and the measured PDFs, despite the relatively small number of ELMs in each sample (a few hundred). We find that inter-burst interval PDFs can be either Gaussian or inverse exponential; furthermore it is possible to distinguish inverse exponential cases that do, or do not, reflect a Poisson process. This suggests that there may be several fundamentally different kinds of underlying processes at work: these include a Gaussian process that has a characteristic waiting time between ELM bursts, as well as a Poisson process with a fixed probability of an ELM per unit time. This provides a new insight into ELM physics, a constraint for future models, and a new statistically rigorous way of quantifying differences between ELM bursts (especially Type III) in different plasmas. The techniques developed here offer a potentially fruitful additional perspective for

ELM physics, which could also be applied in future to ELM bursts at higher frequency and having smaller amplitude.

## Acknowledgments

This work was supported in part by Euratom, the UK DTI, and PPARC.

## References

[1] Connor J W 1998 *Plasma Phys. Control. Fusion* **40** 531
[2] Zohm H 1996 *Plasma Phys. Control. Fusion* **38** 105
[3] Hugill J 2000 *Plasma Phys. Control. Fusion* **42** R75
[4] Hirsch M *et al* 1998 *Plasma Phys. Control. Fusion* **40** 635
[5] Fishpool G M 1998 *Nucl. Fusion* **38** 1373
[6] Zhang W, Tubbing B J D and Ward D J 1998 *Plasma Phys. Control. Fusion* **40** 335
[7] Gill R D *et al* 1998 *Nucl. Fusion* **38** 1461
[8] Connor J W, Hastie R J, Wilson H R and Miller R L 1998 *Phys. Plasmas* **5** 2687
[9] Chapman S C, Dendy R O and Hnat B 2001 *Phys. Rev. Lett.* **86** 2814
[10] Chapman S C 2000 *Phys. Rev. E* **62** 1905
[11] Sornette D 2000 *Critical Phenomena in Natural Sciences* (Berlin: Springer) p 322
[12] Bak P E, Yoshino R, Asakura N and Nakano T 1999 *Phys. Rev. Lett.* **83** 1339
[13] Degeling A W, Martin Y R, Bak P E, Lister Y B and Llobet X 2001 *Plasma Phys. Control. Fusion* **43** 1671
[14] Wegend A S and Gershenfeld N A 1994 *Time Series Prediction* (Reading, MA: Addison Wesley)
[15] Saibene G *et al* 1999 *Nucl. Fusion* **39** 1133
[16] Voss R F 1985 in *Scaling Phenomena in Disordered Systems* ed R Pynn and A Skjeltorp (New York: Plenum) pp 1–11
[17] Greenhough J, Chapman S C, Chaty S, Dendy R O and Rowlands G 2002 *Astron. Astrophys.* **385** 693
[18] Percival D B and Walden A T 1993 *Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques* (Cambridge: Cambridge University Press) pp 111, 374
[19] Wooldridge J M 2002 *Introductory Econometrics: A Modern Approach* 2nd ed (Mason: South-Western College Publishing)
[20] Bak P 1997 *How Nature Works: The Science of Self-Organised Criticality* (Oxford: Oxford University Press)