

An automated sawtooth detection algorithm for strongly varying plasma conditions and crash characteristics

A.Gude, M.Maraschek, O.Kardaun, the ASDEX Upgrade team

¹*Max Planck Institute for Plasma Physics, Garching 85748, Germany*

A sawtooth crash algorithm that can automatically detect irregular sawteeth with strongly varying crash characteristics, including inverted crashes with central signal increase, has been developed. Such sawtooth behaviour is observed in ASDEX Upgrade with its tungsten wall, especially in phases with central ECRH. This application of ECRH for preventing impurity accumulation is envisaged also for ITER. The detection consists of three steps: a sensitive edge detection, a multichannel combination to increase detection performance, and a profile analysis that tests generic sawtooth crash features. The effect of detection parameters on the edge detection results has been investigated using synthetic signals and tested in an application to ASDEX Upgrade soft X-ray data.

1. INTRODUCTION

Sawtooth crashes are a common feature in tokamaks as soon as the central safety factor, q_0 , is below 1 (see [1] and references therein). They are usually preceded by an internal (1,1) kink mode (where poloidal and toroidal mode number are 1). Sawteeth consist of a ramp-phase during which temperatures and particle densities evolve slowly, and a fast crash that basically flattens these quantities inside the mixing radius¹, ρ_{mix} . The position where the values remain constant is the inversion radius, ρ_{inv} . The sawtooth crash decreases the values for $\rho_p < \rho_{inv}$ and increases them for $\rho_{inv} < \rho_p < \rho_{mix}$, provided the pre-crash profile is peaked.

Sawtooth crashes decrease the core confinement but they can help to avoid impurity accumulation in the hot plasma centre, which might be needed to get rid of helium ash in a future fusion reactor [2]. On the other hand, sawtooth crashes can trigger neoclassical tearing modes (NTMs), and the probability for triggering increases for long sawtooth periods [3]. Since fast α particles are expected to have a strongly stabilizing effect on sawteeth [4], future fusion reactors are expected to have very long sawteeth and thus high probability for NTM triggering. Sawtooth control schemes therefore mainly aim at keeping sawteeth small and frequent.

Sawtooth control requires reliable crash detection, which ideally should work for all plasma

¹ Radial coordinates in this paper are all given in the poloidal flux radius, $\rho_p = \sqrt{(\Psi - \Psi_a)/(\Psi_s - \Psi_a)}$, where Ψ is the poloidal flux and the indices a and s refer to the magnetic axis and the separatrix, respectively.




FIG. 1: Time traces of some SXR channels (sample frequency, $f_s = 5$ kHz). Manually detected sawtooth crashes are indicated by vertical dashed lines. Edges with decreasing and increasing signal are enclosed (in selected channels) by dotted triangles pointing down and up, respectively.

scenarios and sawtooth types. Previous sawtooth detection methods use ECE (Electron cyclotron emission radiation) (e.g., [5–8]) or SXR (soft X-ray) (e.g., [9]) measurements. ECE provides a local measurement for the electron temperature, T_e , which is easy to interpret, but ECE measurements can fail when n_e is close to or above the cut-off density [10]. The emissivity in the SXR range depends on T_e and on the densities of electrons and impurities. The T_e dependence is highly nonlinear and varies strongly with the atomic number of the impurities. In addition, SXR measurements are line-integrated. Despite the more difficult interpretation, SXR is well suited for sawtooth crash detection because of its usually good plasma spatial coverage and high availability. In ASDEX Upgrade there are seven SXR pinhole cameras viewing the plasma in the same poloidal plane with a total of 180 lines-of-sight. The sample frequency is 500 kHz or higher and thus more than sufficient to detect sawtooth crashes.

In ITER, metal walls and a tungsten (W) divertor will lead to the presence of heavy impurities in the plasma. In ASDEX Upgrade, with its full W wall, high performance plasmas can usually only be achieved when central wave heating like ECRH is applied to prevent core W accumulation. Under these conditions, SXR profiles are often hollow, i.e. they have a local minimum in the vicinity of the magnetic axis. This is due to hollow tungsten density, n_W , profiles. A sawtooth crash then leads to increase of n_W on axis and decrease further out [11, 12], referred to as inverted crashes in this paper. Often, inverted and normal crashes appear alternatingly (figure 1). Since T_e profiles are usually still peaked when SXR profiles are hollow, there are competing effects: a sawtooth crash increases the central SXR emissivity due to the increase in n_W but decreases it due to the T_e decrease - and vice versa in regions further out. Combined with the non-linear dependence of SXR emissivity on T_e , this leads to a large variety of crash characteristics and often

small crash amplitudes for inverted crashes.

At ASDEX Upgrade ([13], figures 1 and 2), it has been shown that in plasmas with central ECRH and hollow n_W profiles, resistive (1,1) MHD instabilities can flatten T_e inside the $q = 1$ surface before the sawtooth crash. This leads to strongly reduced crash amplitudes. The changes in the central ECE signal can be of the same magnitude as such induced by ELMs (Edge Localized Modes)². Furthermore, a strong variation of ρ_{inv} from crash to crash, depending on the electron temperature precrash profiles, has been observed [13]. Such currently non-standard sawtooth features are likely to be important also in future fusion devices with metal walls. For this reason, we developed a sawtooth crash detector with no prejudice on central crash direction, ρ_{inv} or temporal evolution, which can detect also appearingly small sawtooth crashes.

In section 2 previous sawtooth crash detectors and requirements for a detector for strongly varying sawteeth are discussed. The sawtooth detection algorithm is described in section 3. Section 4 presents a simulation of the influence of detection parameters on edge detection. From this, expectations for detector performance are deduced and compared with results in an application to ASDEX Upgrade SXR data in section 5, followed by a summary in section 6.

2. SAWTOOTH CRASH DETECTORS AND REQUIREMENTS

The most characteristic feature of a sawtooth crash is the abrupt flattening of profiles. Previous detection algorithms therefore search for fast changes (edges) in individual signals, e.g., [5, 7–9], or for an abrupt change towards a centrally flat profile [6]. In TCV [9], a simple difference filter is applied to averaged central SXR signals with thresholds slowly adapted to previous sawteeth. In JET [8] and Tore supra [7], the redistribution character of the sawtooth crash is considered by using difference filtered signals of channels both inside and outside the assumed ρ_{inv} with fixed detection thresholds.

A single-channel edge detection using multi-scale wavelets was applied to ECE data at TEXTOR [5], where the detection reliability was increased by combining several channels. This detection is designed for fast signal increase and decrease similarly and the threshold is self-adaptive. However, in order to keep the detection latency small, each edge direction is predefined, based on previous edges, before the crash detection. The idea in this publication is that ρ_{inv} is close to the $q = 1$ flux surface, which is expected to evolve slowly. Removing this restriction would slow down the

² The effect of ELMs on ECE measurement up to the plasma core is probably similar to the ‘ECE black-out’ reported for JET [14].

detection significantly. Additionally, the method is a plain edge detection and does not include any discrimination for signal edges not caused by a sawtooth crash.

A model-based approach is applied to ASDEX Upgrade ECE data [6]. The T_e profile is compared in steps of 1 ms with model predictions. One model is ‘no crash’, the other models refer to crashes with different ρ_{mix} . This method considers the pre-crash profile explicitly, which should be well suited to distinguish sawtooth crashes from other fast events. However, while the difference methods are in principle applicable to any signal that shows abrupt changes for a sawtooth crash, the model-approach can only compare implemented models, which are currently fully flattened central T_e profiles and ‘no profile change’. Because of the line integration and the unknown density profiles of tungsten (and other impurities), it is doubtful if an implementation based on radiation signals is feasible.

For reliably and automatically detecting sawtooth crashes in a tokamak with metal wall coverage in all scenarios, an algorithm is needed that can cope with irregular sawteeth (in time and amplitude) and with changes in ρ_{inv} (and even in direction of signal edges) from crash to crash. Detection thresholds have to be adapted without any user-intervention. In order to cover also scenarios with densities that are too high for ECE, we use SXR measurements. To serve its purpose, some specific requirements must be fulfilled by the algorithm. Since the signal changes can be small, as discussed above, the detection must be very sensitive. One cannot predict in which plasma region the crashes show up clearly, thus, analysis of signals observing the plasma in different regions is essential. Including signals outside (the usual) ρ_{inv} requires edge detection in both directions, which is necessary already for central SXR channels in case of inverted crashes. Finally, the possible rapid variation of crash characteristics prohibits predefinition of the edge direction, even from the immediate past.

When designing a detector, there is always a trade-off between false negatives (here: missed sawtooth crashes) and false positives (here: detections not caused by a sawtooth crash). A highly sensitive detection in a broad range of signals will lead to a number of false positives, either from real signal edges caused by events that are not sawtooth crashes, or from ‘noise’, including coherent oscillations like the (1,1) mode. To eliminate all false positives effectively, a further discriminating method that considers typical sawtooth crash properties, which apply also for inverted crashes, is required.

3. SAWTOOTH CRASH DETECTION ALGORITHM

To fulfill the above requirements, we designed a sawtooth detection method in three steps:

1. The edge detection is performed by convolution with a suitable kernel, transforming the signal edge into a peak in the convoluted signal. This peak is evaluated and compared to a threshold to find significant edges.
2. Clustering of detections: An edge must be detected in more than one signal within a certain time interval in order to be accepted as possible sawtooth crash.
3. The resulting spatial distribution of signal edges is evaluated versus the expectations for a sawtooth crash. This analysis is carried out only at the preselected time points that passed the first two steps.

In all three steps, design decisions must be made. In this paper we focus on the first step, the optimization of edge detection parameters depending on crash characteristics.

3.1. Convolution with first derivative of Gaussian

In image processing, the Canny algorithm [15] is a common method to extract edges from an image. We use a convolution filter function, $k(\tau)$, proportional to the first derivative of a Gaussian with width σ_k , which is, according to [15], a good approximation for the optimal filter function in terms of reliable detection of edges, precise localization and minimizing multiple responses to a single edge created by noise:

$$k(\tau) := -\tau \frac{1}{\sigma_k^2} e^{-\frac{\tau^2}{2\sigma_k^2}}. \quad (1)$$

For an ‘ideal’ edge signal (without noise), $stp(t)$, which is constant apart from a jump from s_0 to $(s_0 - h_{edge})$ at t_0 ,

$$stp(t) := s_0 - \frac{h_{edge}}{2} (\text{sgn}(t - t_0) + 1), \quad (2)$$

it can be seen that the convolution with $k(\tau)$ gives a Gaussian shaped peak with width σ_k and its extremum equal to the step height, h_{edge} , of $stp(t)$:

$$c(t) := \int_{-\infty}^{\infty} k(\tau) stp(t - \tau) d\tau = h_{edge} e^{-\frac{(t-t_0)^2}{2\sigma_k^2}} \quad (3)$$




FIG. 2: (a)-(c) Time traces for a fast crash, a slow crash and a fast crash with a ramp, respectively. In (c) noise is overlaid. (d) Convoluted signals of all time traces on the calculation frame, l_f . The kernel and its dimensions are depicted in (a) (rescaled in height, not in time), the crash parameters (h_{edge} , w_{edge} , τ_{SC}) in (a)-(c), and l_f in (d).

Owing to the differential filter effect of the kernel, the result depends only on the signal change, not on the signal level, s_0 . The kernel function as defined above has infinite length. Since the kernel edges do not contribute significantly, $k(\tau)$ can be restricted. For practical application we chose $\pm 3.5\sigma_k$ as kernel limits, where the kernel value has dropped to 1.25% of its maximum, such that the kernel width is $l_k = 7\sigma_k$. In reality, we have discrete data, which are digitized with the sample frequency, f_s . Thus, the integrals are replaced by sums over a finite number of data points. Filtering with the kernel as described by equation (1) is analogue to wavelet-based methods, which are often used in image processing and in particular also for sawtooth crash detection [5].

3.2. Peak evaluation

A time trace containing a sawtooth crash, and thus a signal edge with height h_{edge} , results in a convoluted signal with a peak at the edge position (crash time point). The kernel is designed such that an edge towards smaller and larger values produces a positive and negative peak, respectively.




FIG. 3: Relation between minimum reliably detectable sawtooth period, $\tau_{SC,det}$, and detector parameters l_f and $l_k := 7\sigma_k$. (a) Signal with repetitive sawteeth, (b) corresponding convoluted signal. The convolution kernel (height rescaled) is shown in the top box in the position for the last point of the calculation frame around the central peak.

The peak height, p_h , is similar to h_{edge} for an abrupt crash, but it decreases with increasing crash duration, w_{edge} (see figure 2)³. However, real signals also comprise noise, which results in local extrema in the convoluted signal. The detection has to distinguish between significant peaks, representing signal edges, and insignificant (noise) peaks. A suitable way is a threshold for the peak height. For automated detection this has to be self-adaptive to the signal, which is fulfilled, e.g., for thresholds based on a noise level.

In order to allow fast variations of sawtooth characteristics, the background reference (noise level) should be determined in the vicinity of the crash. Therefore, we analyse the signal in a calculation frame with length l_f , spreading symmetrically around the peak. For the algorithm to work reliably, we must assume that only one sawtooth crash is within this calculation frame, so that the corresponding peak is the largest local and thus the global extremum within its frame. For a given smallest detectable sawtooth period, $\tau_{SC,exp}$, this condition implies an upper limit for $l_f + 7\sigma_k$ (see figure 3):

$$\tau_{SC,det} := \frac{1}{2}(l_f + 7\sigma_k) \leq \tau_{SC,exp}. \quad (4)$$

In the following we describe three different methods for peak evaluation, two of them making use of the variance of the convoluted signal for threshold determination, the third one using the kurtosis instead. All methods are unprejudiced with respect to the direction of the signal edge.

³ For $w_{edge} = \sigma_k$: $p_h \approx 0.94h_{edge}$, for $w_{edge} = 7\sigma_k$: $p_h \approx 0.36h_{edge}$.

3.2.1. POS - Peak-over-Sigma

A threshold, $p_{h,min}$, for the absolute peak height, $|p_h|$, can be based on the noise level of the convoluted signal, $c = c_i$ ($i \in \{1, \dots, n\}$, where n is the number of data points in the calculation frame). The sample estimate of the standard deviation, σ_c , is

$$\hat{\sigma}_c = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (c_i - \hat{\mu}_c)^2}, \quad (5)$$

with $\hat{\mu}_c$ the estimated mean of c . Here we use $\hat{\sigma}_c$ as a practical measure of the noise level of c . Owing to the convolution with the kernel, σ_c (and also $\hat{\sigma}_c$) for a pure noise signal is different from the standard deviation of the original noise signal, σ_s . In [15], section VI, the autocovariance function of white noise filtered by the first derivative of a Gaussian is given. This includes an unspecified proportionality factor (k_3), the absolute value of which is equal to the variance, σ_c^2 , of the convoluted signal, c . For the kernel in equation (1) one can derive $\sigma_c^2 = \sqrt{\pi} \sigma_s^2 / (2\sigma_k f_s)$, where $\sigma_k f_s$ corresponds to the number of points within σ_k . For $\sigma_k f_s = 14.286$ (as used in the simulation in section 4) and $\sigma_k f_s = 2.0$ (which is - according to our experience with ASDEX Upgrade time traces - a safe lower limit for sufficient kernel resolution), one gets $\sigma_c \approx 0.25\sigma_s$ and $\sigma_c \approx 0.66\sigma_s$, respectively.

An adaptive threshold is defined by $p_{h,min} := th_{POS} \cdot \hat{\sigma}_c$, where th_{POS} is a constant detector parameter, which only depends on the detector settings, but not on the edges to be detected. Thus, a possible condition for peak acceptance is $|p_h|/\hat{\sigma}_c \geq th_{POS}$, where p_h is the (signed) c_i with the highest absolute value. However, a ramp with constant inclination in the signal leads to an offset of c (see figure 2, (c) and (d)). Since in sawtooth signals, ramps are a common feature, the mean of c has to be subtracted from the peak height and thus we define the *peak weight* for the POS-method as

$$W_{POS} := \frac{|p_h - \hat{\mu}_c|}{\hat{\sigma}_c}. \quad (6)$$

Peaks with $W_{POS} \geq th_{POS}$ are accepted.

Here, $\hat{\mu}_c$ and $\hat{\sigma}_c$ are calculated from the entire signal, so the peak itself contributes to these values, which leads unavoidably to an underestimation of the peak weight, especially for narrow calculation frames.

3.2.2. POSR - Peak-over-Sigma-reduced

In order to avoid the influence of the peak on mean and standard deviation, the peak can be excluded from their calculation. The peak does not only consist of one data point, but has a finite width and thus several points increased with respect to the background. Its width is proportional to the kernel width, σ_k , as long as the edge is not too broad, i.e., the crash is not too slow. We exclude the number of data points corresponding to the full kernel length, $l_k = 7\sigma_k$, to get $\hat{\mu}_r$ and $\hat{\sigma}_r$, the estimated mean and standard deviation of the reduced data set, respectively. However, instead of excluding the data points inside $\pm 3.5\sigma_k$ around the position of the peak extremum, we exclude the largest absolute values within the whole calculation frame, no matter at which time they appear. This results in a better noise reference, and thus more reliable edge detection, in the presence of a second (slightly smaller, negative or positive) peak in the calculation frame. A similar method has been applied in [15]. The peak weight is defined as

$$W_{POSR} := \frac{|p_h - \hat{\mu}_r|}{\hat{\sigma}_r}. \quad (7)$$

3.2.3. Kurtosis

The third method we have applied makes use of the kurtosis, of which the large sample analogue is

$$Kurt(c) := \frac{1}{n} \sum_{i=1}^n \left(\frac{c_i - \hat{\mu}_c}{\hat{\sigma}_c} \right)^4. \quad (8)$$

We use the kurtosis as a measure of the significance of outliers, since in view of the 4th power, they strongly influence $Kurt(c)$. Applied to the convoluted signal, it is thus a measure for the significance of the peak with respect to the signal variance. Therefore equation (8) is applied to the whole data set in a calculation frame, also in cases of relatively small n , where the sample variance of this estimator is quite large. The peak weight for the kurtosis method is simply defined as

$$W_{Kurt} := Kurt(c). \quad (9)$$

A second peak within the calculation frame increases W_{Kurt} and thus the sensitivity of the kurtosis method - as long as there is also sufficient ‘normal’ background in the frame.

3.3. Multichannel analysis

Using multiple signals increases the detector performance by increasing the probability of detecting every sawtooth crash. Multichannel analysis can also help in reducing the false positive rate, as will be discussed below.

3.3.1. Clustering

Detection due to pure noise is unlikely to be observed simultaneously in several signals. The condition that an edge must be observed in at least two channels is therefore already effective in reducing false positives. A sawtooth crash appears (nearly) simultaneously in all channels inside the mixing radius. However, detections in different signals for the same crash will - with high temporal resolution - not always be observed at the very same sample. Therefore, a time interval, t_{clust} , must be defined, during which detections are regarded as belonging to the same edge. This should be chosen as short as possible in order not to combine detections from different events, but also long enough to catch all detections from the same event. We use $t_{clust} = 0.5$ ms, which is above the typical time scatter with which crashes show up in different areas of the mixing region, but sufficient to separate edges as close as 1 ms.

3.3.2. Profile analysis

The line-integrated SXR signals are mapped to the minimum ρ_p along the line-of-sight. To distinguish between the two sides of the camera with respect to the magnetic axis, the signals on one side are labelled with negative ρ_p values. In the following, the difference of the SXR profiles, observed from a single SXR pinhole camera, before and immediately after a detected edge, is referred to as crash profile, irrespective of the type of detected event.

The crash profile has characteristic features in case of a sawtooth crash. However, the classical assumptions - a sawtooth crash reduces T_e and radiation in the centre, and ρ_{inv} and the sawtooth period are slowly evolving quantities - are no longer applicable for sawtooth characteristics as described in section 1. We therefore use only history-independent assumptions that test generic sawtooth features. Although ρ_{inv} is not necessarily identical to the $q = 1$ location [16] and has been observed to vary strongly from crash to crash, the *existence* of an inversion is an important feature of a sawtooth crash. Furthermore, a sawtooth crash profile must generally be roughly symmetric with respect to the magnetic axis. In [13] we have already presented a sawtooth discrimination




FIG. 4: (a) Time trace of a central SXR signal (black) with the kernel (white curve) in arbitrary units, (b) profiles of SXR signals averaged over the time ranges marked in (a), (c) difference of profiles in (b), (d) profile of convoluted signals at crash time point. Loss areas (A^-) and gain areas (A^+) are filled in (c) and (d).

method for SXR signals using empirical parameters, which we repeat here for convenience.

The crash profile is represented well by the convolutions of the signals observed at different ρ_p , taken at the edge time point, $c(\rho_p)|_{t_{SC}}$ (see figure 4, (c) and (d)). We chose a region $|\rho_p| \leq \rho_m := 0.7$ in order to safely cover the full mixing region for the profile analysis. The channel selection for edge detection was smaller but also included channels outside ρ_{inv} . The crash profiles, originally only defined at the SXR measurement positions, are interpolated to an equidistant grid in ρ_p for the following calculations.

The first parameter represents the significance of a crash event:

$$A_{norm} := \frac{A_{crash}}{A_{rad}} = \frac{A^+ + A^-}{A_{rad}} = \frac{\sum (|c(\rho_p)| + |c(-\rho_p)|)}{\sum (s(\rho_p) + s(-\rho_p))}, \quad (10)$$

where c is the convoluted and s the original signal at t_{SC} . This and the following summations are over the n_p interpolated values with $0 < \rho_p \leq \rho_m$. A^+ and A^- represent the areas of gains and losses in the crash profile, respectively (figure 4). A_{norm} might serve as crash amplitude, but it strongly depends on the precrash profile shape, and a proper quantification is still to be determined. We found values as low as 2 % for inverted crashes in SXR, which is due to the competing effects of W density and T_e redistribution. Thus, we set the condition for acceptance as sawtooth crash

to $A_{norm} \geq 2\%$.

The second parameter is determined by the redistribution:

$$A_{net} := \frac{|A^+ - A^-|}{A_{crash}} = \frac{|\sum (c(\rho_p) + c(-\rho_p))|}{A_{crash}} \quad (11)$$

is the absolute value of the integrated crash profile divided by A_{crash} . We set the criterion for acceptance as $A_{net} < 90\%$, i.e., neither A^+ nor A^- must be smaller than 5% of the total crash area, $A_{crash} = A^+ + A^-$.

One type of false positives arises from slow strong (1,1) mode oscillations with large growth or decay rate. Since the (1,1) mode has opposite phase on opposite sides of the magnetic axis, the related crash profiles are strongly asymmetric for mode induced detections. A similar effect can appear for ELMs: they can lead to displacements and distortions of the plasma column that cause the SXR line-integrated emissivity to increase on one side of the magnetic axis and decrease on the other side. The third parameter therefore quantifies the profile asymmetry:

$$p_{asym} := \frac{\sqrt{\sum (c(\rho_p) - c(-\rho_p))^2}}{A_{crash}}. \quad (12)$$

This parameter can only be applied to measurements that extend to both sides of the magnetic axis, such as SXR at ASDEX Upgrade. We set the criterion for accepted sawtooth crashes to $p_{asym} \leq 1$. A drawback is that also sawtooth crashes that are preceded by a very slow (1,1) mode can be rejected, depending on the mode phase at t_{SC} .

The threshold values of the profile criteria are determined empirically from analysis of a few ASDEX Upgrade discharges. This has to be refined in the future, e.g., by replacing the fixed numbers by a combination of probabilities and by consideration of precrash profile features. In this paper we concentrate on the edge detection. However, we will keep in mind that the analysed crash profiles are the result of the convolution. Thus, the convolution kernel should be chosen such that the crash profiles represent the sawtooth crash well.

4. EFFECT OF DETECTOR PARAMETERS ON SYNTHETIC SIGNALS

The first step in our sawtooth detection algorithm is the edge detection in individual channels. Once the peak evaluation method and kernel shape are chosen, there are still three parameters to optimize for an edge detection based on convolution with a kernel as defined in equation (1): The width of the convolution kernel, σ_k , the frame in which the peak weights are calculated, l_f , and the threshold th_j ($j \in \{\text{POS, POSR, Kurt}\}$) for the peak weight, W_j . The detection parameters




FIG. 5: Peak weight thresholds, $\langle th_j \rangle$, for the three methods ($j \in \{\text{POS}, \text{POSR}, \text{Kurt}\}$) determined from pure noise signals without edge versus normalized frame length, $\|l_f\|$, for normalized sample frequency $\|f_s\| = 14.286$ and various acceptable false positive rates, fpr_{acc} , (from 0.15 % (darkest) to 5 % (lightest)).

have to be adjusted to the boundary conditions. These are determined by the sawtooth crashes to be detected: the minimum sawtooth period, $\tau_{SC,exp}$, the maximum edge width, $w_{edge,max}$ (crash duration) and the minimum edge height, $h_{edge,min}$ (signal change due to crash). Furthermore, the signal-to-noise-ratio and the sample rate, f_s , of the diagnostic signal as well as the quality standards (acceptable false positive rate, fpr_{acc} , and acceptable false negative rate, fnr_{acc}) have to be considered. Another important influence on detection results comes from coherent perturbations, such as the sawtooth precursor. This is not included in the simulation but discussed later as an additional constraint.

To study the influence of parameter variation, we produce normally distributed noise, with mean zero and standard deviation, σ_s , by a random number generator. To these noise signals we add an edge symmetrically around $t = t_{SC}$, which can be varied in its width, w_{edge} , and its height, h_{edge} . Pure noise signals are represented by $h_{edge} = 0$. To account for the repetitive sawtooth character, a ramp between sawtooth crashes can be added. The sawtooth period is the sum of ramp and crash duration (see figure 2). For each parameter set, we generate $N = 4000$ signals, which are mutually uncorrelated in their noise component, and calculate the peak weights, W_j (with $j \in \{\text{POS}, \text{POSR}, \text{Kurt}\}$). From these distributions with 4000 values, the error rate for a given threshold can be determined with a resolution of 0.025 %. For a signal with an edge, an error (a false negative) is a peak weight below the threshold, while for pure noise signals, a peak weight above the threshold is an error (a false positive). Additionally, we have generated one noiseless signal for each parameter set to determine the noiseless detection limits.

4.1. Threshold determination

Before detecting edges in the signals, a reasonable threshold for the peak weight has to be determined, which has to reject *most* peaks that result from noise. For this purpose, we generate signals with no edge but only noise. The parameters for noise signals are: sample rate, f_s , peak weight method, frame length, l_f , kernel width, σ_k , and the acceptable false positive rate, fpr_{acc} . Since scaling in the temporal direction does not modify the data points, there are only two parameters independent out of f_s , l_f and σ_k . We normalize the temporal values to σ_k , i.e. $\|l_f\| := l_f/\sigma_k$ and $\|f_s\| := f_s\sigma_k$.

For a given threshold, th_j , all $W_j \geq th_j$ are false positives. So - within the same distribution - each threshold value is related to a false positive rate, fpr , through the complementary cumulative distribution function, $th_j \mapsto fpr(th_j)$ with domain \mathbb{R}^+ . However, we have to find the appropriate th_j for a given fpr_{acc} . Therefore, we are interested in the inverse function, $th_j(fpr_{acc})$, which, strictly speaking, does not exist for discrete data. However, for $N = 4000$, it can be approximated with sufficient precision by the $100(1 - fpr_{acc})$ th percentile of the distribution⁴.

In order to reach a small uncertainty of the threshold values, which are important for all our further studies, we generate $M = 50$ distributions (each consisting of $N = 4000$ signals) for each parameter set. For each fpr_{acc} value ($1/4000 \dots 3999/4000$) we calculate the mean, $\langle th_j \rangle$, and the empirical standard deviation, $\hat{\sigma}_{th_j}$, of the $M = 50$ th_j values. The uncertainty of $\langle th_j \rangle$ is given by the estimated standard deviation of the mean, $u = \hat{\sigma}_{th_j}/\sqrt{M}$. The simulation was performed only for $\|l_f\| \geq 8$, in order to leave a sufficient part of the data set for the calculation of $\hat{\sigma}_r$ and $\hat{\mu}_r$.

Figure 5 shows the dependencies of the thresholds and their uncertainties, u , on $\|l_f\|$ for our three methods. The most relevant range for application is $10 < \|l_f\| < 60$, as will be discussed in section 5.

For POSR there is a strong increase of the threshold for small $\|l_f\|$. This results from the elimination of the largest absolute values, which leads to reduction of $\hat{\sigma}_c$ for small $\|l_f\|$, and thus increase of W_{POSR} . In the investigated range $0.15\% \leq fpr_{acc} \leq 5\%$, fpr_{acc} has a relatively small effect on the thresholds. However, its effect on detection still has to be determined. The dependency on $\|f_s\|$ is very weak (not shown). The thresholds do not depend on the noise level itself, because of the scaling invariance of the peak weights. In the following simulations and also for the application to ASDEX Upgrade data, described in section 5.5.3, $\langle th_j \rangle$ (for the corresponding

⁴ The peak weights are counted, starting from the highest, until the number of false positives corresponding to fpr_{acc} is reached. The corresponding value of W_j is taken as th_j for the chosen fpr_{acc} .

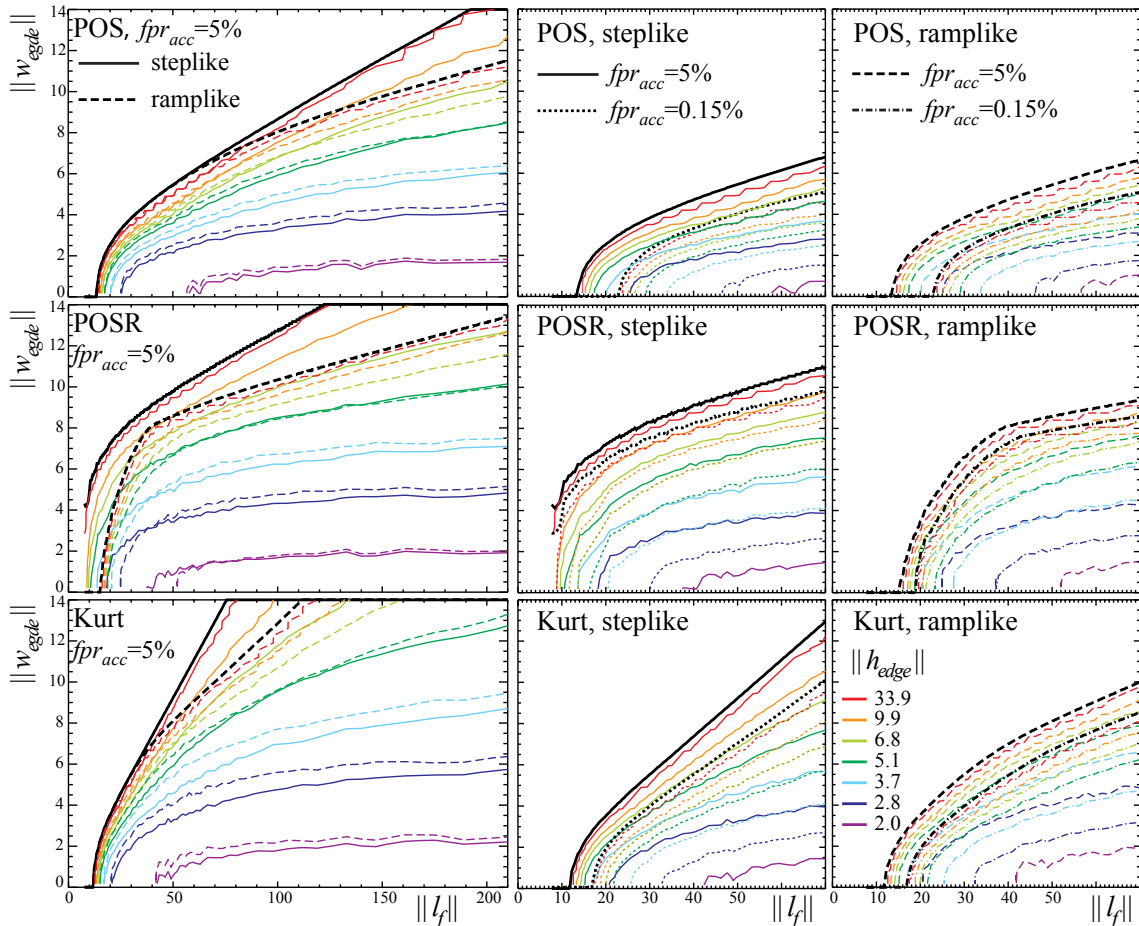


FIG. 6: Contour lines for ‘reliable’ detection ($fnr \leq 0.15\%$) for the three methods: POS (top), POSR (middle), $Kurt$ (bottom) with normalized sample frequency $\|f_s\| = 14.286$. The noiseless limits are plotted as bold black lines. The left column compares steplike (solid lines) and ramplike (dashed lines) signals for $fpr_{acc} = 5\%$. The middle and right column repeat these limits for $fpr_{acc} = 5\%$ in a smaller range of $\|l_f\|$ and compare them to limits for $fpr_{acc} = 0.15\%$ for both signal types (middle: steplike, right: ramplike signals). The legends for the line styles in the top boxes apply to the respective columns. The colour code is shown in the bottom right box.

$\|l_f\|$, $\|f_s\|$, method j and fpr_{acc}) serves as peak weight threshold.

4.2. Steplike signals

Simulation of time traces containing a sawtooth crash requires additional parameters: the edge height, h_{edge} , and its width, w_{edge} , which is the crash duration. First, we consider time traces which are at a constant level outside the crash. Since the signal level itself is not visible in the convoluted signal, the relative signal change, h_{edge}/s_0 , where s_0 is the signal before the edge, is not

a suitable simulation parameter. We use $\|h_{edge}\| := h_{edge}/\sigma_s$ instead, where σ_s is the noise level, represented by the standard deviation, of $s(t)$. Similar to the other temporal parameters, the edge width is normalized to the kernel width, $\|w_{edge}\| := w_{edge}/\sigma_k$. A typical value for the sawtooth crash duration in ASDEX Upgrade is 0.1 ms, but ≈ 1 ms is occasionally observed.

We generate $M = 5$ distributions at a fixed normalized sample rate of $\|f_s\| = 14.286$, corresponding to 101 data points within $l_k = 7\sigma_k$. From these signals, the peak weights are calculated for our three methods and compared with the respective threshold. Since there is an edge in the signals, which should be detected, now all values *below* the threshold are errors: false negatives. The rate of false negatives depends on the kernel width and the frame length, and additionally on fpr_{acc} , which determines the thresholds. The variances of the false negative rates are ignored in the following.

In figure 6 (left column) the limits for ‘reliable’ detection with $fnr_{acc} = 0.15\%$, and thus high sensitivity, are shown for all three methods as solid contour lines in the $\|l_f\|$ - $\|w_{edge}\|$ -plane for various $\|h_{edge}\|$ and $fpr_{acc} = 5\%$. A particular edge with $\|w_{edge}\|$ below and above the contour line for its $\|h_{edge}\|$ will be detected with a probability above and below 99.85%, respectively. Large $\|w_{edge}\|$ and small $\|l_f\|$ are detrimental for edge detection. For large $\|h_{edge}\|$, the contour lines approach the one for noiseless signals, marked by a bold black line. The most remarkable feature is that the POSR method works reliably with significantly smaller $\|l_f\|$ than the other two methods.

The middle column of figure 6 depicts the effect of fpr_{acc} on detection. The limits are shifted to larger $\|l_f\|$ for smaller fpr_{acc} . Small edges are more affected than large ones. The shift for large edges is smallest for the POSR and largest for the POS method.

4.3. Ramplike signals

Sawteeth are repetitive events, so we have to consider the neighbouring crashes. The detection of a peak in c is reliable only if there are no other peaks (of similar or larger size) in the same calculation frame, which is reflected by equation (4) and figure 3. In a stationary plasma with regular sawteeth, the signal after a crash recovers to its precrash value before the next crash. There is an infinite number of possible signal shapes between the crashes. However, a typical one, which led to the name sawteeth, is a ramp with constant inclination. The biggest effect is expected for the steepest ramp, which is correlated to $\tau_{SC,exp} = 1/2(l_f + 7\sigma_k)$ (see figure 2, (c)). In order not to add another crash parameter to the simulation, we restrict on analysing signals with this shortest possible ramp with constant inclination. Thus, we analyze only the extreme signal

shapes: fully steplike and fully ramplike signals. Slower ramps should lead to detection results in between the two extrema. However, strongly non-stationary variations of the signal level can lead to reduction of the peak weight, mainly due to increased $\hat{\sigma}_c$, and thus possible rejection of an edge.

The limits for ‘reliable’ detection of a sawtooth crash surrounded by ramps with constant inclination are shown in figure 6, left column, as dashed lines for $fpr_{acc} = 5\%$. For all methods, the tolerable edge width at large $\|h_{edge}\|$ and large $\|l_f\|$ is smaller than for steplike signals. Only for the POSR method there is also a strong deviation for small $\|l_f\|$: its detection capability at very small $\|l_f\| \leq 15$ vanishes for ramplike signals. The ramp character relevant for the detection depends on the detector settings: for $\tau_{SC,det} \ll \tau_{SC,exp}$ the analyzed signal part covers a small fraction of the sawtooth period and thus the signal is almost steplike.

The effect of fpr_{acc} is shown in the right column of figure 6 and is very similar to that observed for steplike signals.

4.4. Sample frequency

When a signal, consisting of pure white noise, is sampled with a lower frequency⁵, f_s , σ_s scales with $\sqrt{f_s}$, such that $\|h_{edge}\|$ scales with $1/\sqrt{f_s}$. The minimum detectable $\|h_{edge}\|$ in the simulations, at fixed detector and other edge parameters, also scales with $1/\sqrt{f_s}$. Thus, the detection of an edge does not depend on the sample rate.

The former is valid if the kernel is the dominant low-pass filter. The absolute value of the Fourier transform of the kernel function (see equation (1)) has similar shape as the kernel with the amplitude maximum at $f_{res} := 1/(2\pi\sigma_k)$. The amplitude at 2 (3, 4) f_{res} is at 45 (5.5, 0.22) % of its maximum, such that frequencies beyond $3f_{res}$ are damped effectively by the kernel. When the highest detectable frequency, $f_N := f_s/2$, is well above this limit, the edge detection result is insensitive to f_s . One point per σ_k interval (i.e. $f_N = \pi f_{res}$) is marginal for effective kernel filtering and also does not allow to span the full kernel region up to $\pm 3.5\sigma_k$. Therefore, it seems sensible to take at least two points per σ_k interval, corresponding to $f_s = 4\pi f_{res} = 2/\sigma_k$.

⁵ The precondition is that the power spectral density up to $f_s/2$ is preserved. This requires a suitable lowpass filter, either included in the downsampling of already digitized data, or as a hardware component for sampling of an analog signal with a smaller frequency.

5. CHOICE OF DETECTOR SETTINGS

The experimentally determined quantity h_{edge}/s_0 can be regarded as the sawtooth crash amplitude. In principle, it is connected with the simulation relevant quantity $\|h_{edge}\| := h_{edge}/\sigma_s$ by the signal-to-noise ratio of the measured signal. However, the ‘noise’ in real time traces usually is not white but has higher amplitudes at characteristic frequencies. This is especially true for sawtooth detection, since a feature of the sawtooth crash is a precursor. Thus, a comparison of the noise levels of measured time traces with the simulation only makes sense for the convoluted signals using the same kernel width. Moreover, fast growth and decay rates of perturbations lead to non-stationary spectra and thus a dependency of the signal variance on the calculation frame length. Considering also the statistical nature of the simulation results, it is not possible to predict the success of edge detection for each individual sawtooth crash. Still, a suitable range of detector settings for given boundary conditions can be derived. This will be shown by comparison of expectations based on the simulation with results of the sawtooth detection for ASDEX Upgrade SXR data.

5.1. Influence of perturbation frequencies

In addition to the simulation results, the influence of perturbation frequencies must be considered. Oscillations can cause problems for the crash detection if their frequency is below $2 \dots 3f_{res}$, since then their amplitude is not sufficiently suppressed by the kernel filtering. In such cases two things can happen. Either the edge detection can fail or the profiles can be rejected due to large asymmetry. The latter is the consequence of the odd (1,1) sawtooth precursor symmetry. In order to avoid the mode influence, the lowest frequency⁶ of significant perturbations, f_{mode} , sets a lower limit to the kernel width:

$$\sigma_k \geq \frac{1 \dots 1.5}{\pi f_{mode}}. \quad (13)$$

Using equation (4), one gets $\|l_f\| = 2\tau_{SC,det}/\sigma_k - 7 = 4\pi\tau_{SC,det}f_{res} - 7$. For the condition $f_{mode} \geq 2f_{res}$ this gives $\|l_f\| \leq 2\pi\tau_{SC,det}f_{mode} - 7$, such that $\|l_f\|$ has an upper limit for a given f_{mode} and $\tau_{SC,det}$ (see figure 7).

⁶ Here we only consider frequencies above f_{res} . Very low frequencies are also damped by the kernel filtering. At 0.28 and 0.033 f_{res} the amplitude damping is similar to that at 2 and 3 f_{res} , respectively.




FIG. 7: Lines of maximal $\|l_f\|$ in the $\tau_{SC,det}$ - f_{mode} plane for the condition $f_{mode} \geq 2f_{res}$. The dots mark the first three settings in table I. The dotted, dashed-dotted and dashed bold lines mark $\|l_f\| = 42$, 21 and 14, respectively.

5.2. Expected effects of parameter choice

The two main detector parameters to choose are the ratio and the sum of l_f and $l_k = 7\sigma_k$, the latter corresponding to $2\tau_{SC,det}$. Considering the simulation only (see figure 6), $\|l_f\|$ for detecting extremely small sawtooth crashes is bounded below to $\|l_f\| > 40$ for POSR and kurtosis, and to $\|l_f\| > 60$ for POS. However, increasing $\|l_f\|$ beyond this should not necessarily improve the detection. Increasing $\|l_f\|$ with constant $\tau_{SC,det}$ results in a shorter kernel and thus higher f_{res} , such that higher perturbation frequencies can disturb the sawtooth detection. Increasing $\|l_f\|$ and also $\tau_{SC,det}$, such that the kernel width is constant, leads to a larger minimum detectable sawtooth period, which can result in missed crashes when the sawtooth period in the signals, $\tau_{SC,exp}$, is smaller than $\tau_{SC,det}$. Additionally, it has to be kept in mind that the simulation considers only ramps with constant inclination (including flat signals) around the crash. For strongly non-stationary signals, we cannot expect to see the same positive effect of increasing $\|l_f\|$.

For a real-time implementation, the latency for a detection of an individual sawtooth crash should be as small as possible. The time range after a sawtooth crash, that is used for detection is equal to $\tau_{SC,det} = (l_f + 7\sigma_k)/2$. Reducing l_f will usually force a reduction of $\|l_f\|$ in order to keep f_{res} low enough. Therefore, POSR is expected to be the best of the three methods for short latency, since it allows very small $\|l_f\|$ for steplike signals. In order to keep this advantage of steplike signals, the sawtooth period in the signals, $\tau_{SC,exp}$, should be significantly larger than $\tau_{SC,det}$, because only then the ramp character of the crashes is small. For fully ramplike signals, the kurtosis method is expected to detect sawteeth at smaller $\|l_f\|$ than the POSR method. However,

with small $\tau_{SC,det}$ it is likely that $\tau_{SC,det} \ll \tau_{SC,exp}$.

name	l_f (ms)	σ_k (ms)	$\tau_{SC,det}$ (ms)	$\ l_f\ $	f_{res} (kHz)	$f_{s,min}$ (kHz)
small crashes	10.3	0.25	6.0	42	0.65	8
short latency	1.33	0.095	1.0	14	1.67	21
intermediate	5.18	0.25	3.45	21	0.65	8
11 ms steplike	5.33	0.38	4.0	14	0.42	5.3
11 ms ramplike	14.7	1.05	11.0	14	0.15	1.9

TABLE I: Detector parameters and correlated quantities for the settings applied in this paper: calculation frame l_f , kernel width σ_k , minimum reliably detectable sawtooth period $\tau_{SC,det} := (l_f + 7\sigma_k)/2$, normalized frame length $\|l_f\| := l_f/\sigma_k$, kernel resonance frequency $f_{res} = 1/(2\pi\sigma_k)$ and recommended minimum sample frequency, $f_{s,min}$.

Table I lists the detector settings used in the application. For the ‘small crashes’ and ‘intermediate’ settings we designed the kernel to damp perturbations with frequencies $f_{mode} \geq 1.25$ kHz by at least a factor of two. The ‘short latency’ setting uses only a short time interval of $\tau_{SC,det} = 1$ ms after the crash for its determination but requires a smaller kernel width with higher f_{res} . Two special settings for testing the effect of varying $\tau_{SC,det}$ in the low $\|l_f\|$ -region in discharge #31258 are also listed in table I.

The crash duration, w_{edge} , and the perturbation frequency, f_{mode} , both cause a lower bound for the kernel width. At ASDEX Upgrade, a typical crash duration is ≤ 0.1 ms, while some crashes last up to around 1 ms. Let us consider the example of the POSR method and steplike sawteeth (see central box in figure 6 or figure 12). For the ‘small crashes’ setting, the maximum ‘reliably’ detectable normalized edge width is $\|w_{edge}\| \approx 8$ for relatively large crashes (large $\|h_{edge}\|$), corresponding to $w_{edge} \approx 2$ ms. Thus, having eliminated the influence of perturbation frequencies with $f_{mode} \geq 1.25$ kHz, a negative effect of the crash duration can mainly be expected for very small sawtooth crashes. For the ‘short latency’ setting, large crashes are limited to $\|w_{edge}\| < 5$, corresponding - due to the smaller kernel width - to $w_{edge} \approx 0.5$ ms. For this setting, slow crashes will hamper the detection even if they are very large.

For the POS method, a second peak in the frame increases the variance of the data set and thus the peak weight is decreased (see equation (6)). For POSR, this effect is strongly reduced by design, because we ignore the data points with the largest absolute values, no matter in which part of the frame they lie. The kurtosis method is intrinsically insensitive to the position of outlier values,

such that a second peak in the detection frame does not decrease the sensitivity. In conclusion, POS should not only have disadvantages for very small crashes, but also for such that are closer to each other than $\tau_{SC,det}$.

	method	fpr_{acc} (%)	setting	f_s (kHz)
trace A	POS	5.0	small crashes	20
trace B	POSR	5.0	small crashes	20
trace C	Kurt	5.0	small crashes	20
trace D	POSR	5.0	short latency	100
trace E	POSR	5.0	intermediate	100
trace F	POSR	0.15	small crashes	100
dots	all	5.0	short latency	100

TABLE II: Detector parameters and sample frequency for the results shown in figure 8.

5.3. Application to ASDEX Upgrade soft X-ray data

To demonstrate the effect of detection parameters for experimental data, we choose discharge #33186, because of its strongly different heating schemes (figure 8, top) and thereby sawtooth phases. In the early phase (1.0-1.4s) with only ohmic heating, the frequency of the precursor, which exists only for a half cycle for many crashes, is about 1.25 kHz. In the phase 1.4-2.0s, when NBI heats the plasma, the plasma rotation (and hence the (1,1) frequency) speeds up. During the ECRH heating phase (2.0-4.0s), hollow SXR emissivity profiles and inverted sawtooth crashes, as described in [11], appear. The sawtooth period is strongly decreased with respect to the pure NBI phase and additionally intermediate crashes appear. In the phase with ICRF heating (4.0-6.0s), SXR profiles are peaked. In the following, we compare the results of the sawtooth detection, including clustering (with $t_{clust} = 0.5$ ms) and profile analysis, with the expectations derived in section 5.5.2.

Figure 8 (middle graph) shows all edge detections after clustering for discharge #33186 from 1.0s to 6.4s with the ‘small crashes’ setting (see table I) for the POSR method and $fpr_{acc} = 5\%$. As expected from the rather high fpr_{acc} there are several false positives after edge detection in the sense that they do not correspond to sawtooth crashes. Especially in the late phase with large long sawteeth ELMs lead to several edge detections.

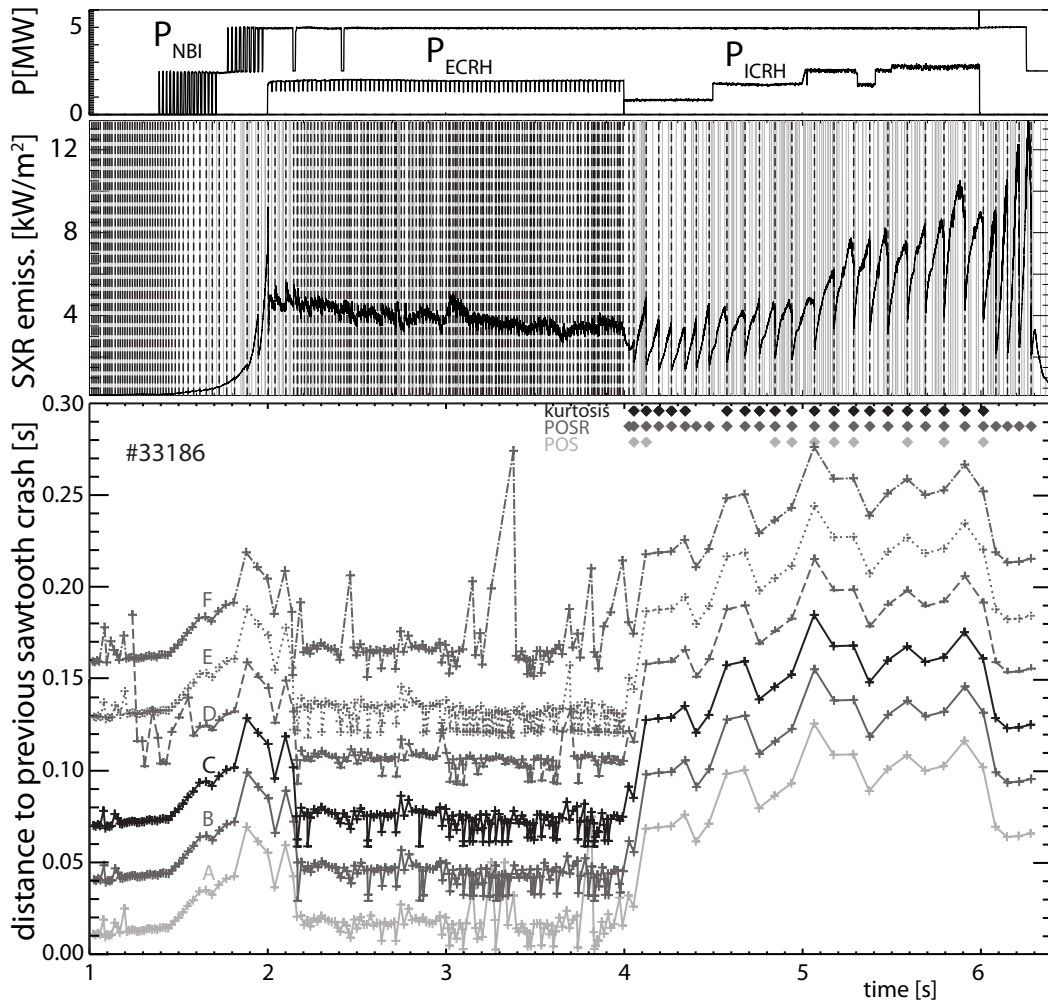


FIG. 8: Top: heating time traces for discharge #33186. Middle: central SXR channel (sample frequency $f_s = 20$ kHz) with clustered edge detections for POSR method with parameters for ‘small crashes’ (see table I) and acceptable false positive rate $fpr_{acc} = 5\%$. Accepted edges after profile analysis are depicted by vertical dashed bold lines, rejected ones by grey narrow lines. Bottom: time difference to the previous detected crash for various detector settings, each trace shifted up by 30 ms with respect to the previous one. Table II lists the applied settings, methods, fpr_{acc} and frequencies for traces A-F and the dots in the top right corner.

In the early ohmically heated phase nearly all crashes are detected by each of the three methods, as can be seen by the smooth sawtooth period traces A-C in figure 8, bottom. In the late ICRF heated phase with big sawteeth all sawtooth crashes are clearly detected and all false positives rejected by the profile analysis. At first sight, the variation in the sawtooth period between 2 s and 4 s, seems to show detection errors, namely not rejected false positives. However, figure 9 reveals that the detected short sawtooth periods are correct. Since $\tau_{SC,det} = 6$ ms was chosen




FIG. 9: Zoom of one central SXR signal and two ECE signals, one inside, the other outside the sawtooth inversion radius, in the phase with inverted crashes (in SXR signals). Dashed vertical lines depict the edge detections with SXR (‘small crash’ setting), dotted vertical lines show the manually detected sawtooth crashes in ECE signals.

and the distance between crashes can be as short as ≈ 3 ms, detection of these close-by crashes is not expected to be complete⁷. The detection results of POSR and kurtosis are nearly identical, the deviations only being in few close-by crashes. POS misses some crashes completely, i.e., not even one of a group of close-by crashes is detected. This is because for two or more peaks in one calculation frame, the POS peak-weight is reduced.

The effect of a small kernel width in the presence of perturbation frequencies is demonstrated by the detection of a single sawtooth crash (discharge #33186 at 6.084s) with $\tau_{SC,det} = 1$ ms and two different $\|l_f\|$ (see figure 10). With $\|l_f\| = 42$, the kernel width is $\sigma_k = 0.041$ ms and $f_{res} = 3.9$ kHz, such that even the high precursor frequencies around 8 kHz in the ICRF phase hamper the sawtooth detection. The crash profile is too asymmetric and rejected by the profile analysis. The longer kernel with $f_{res} = 1.67$ kHz for $\|l_f\| = 14$ (which corresponds to the ‘short latency’ setting) results in a sufficiently symmetric crash profile. Scanning f_{res} , the sawtooth crash is accepted up to 3.3 kHz. This corresponds to $f_{mode} \approx 2.4f_{res}$, which is in the expected range for the disappearing influence of a large mode. The detection results for the POSR method with the ‘short latency’ setting ($\|l_f\| = 14$) in discharge #33186 are represented by trace D in figure 8. In the early phase with small perturbation frequencies, this setting is not suitable, while in the late phase sawtooth crashes are detected reliably. The superiority of the POSR method at $\|l_f\| = 14$ for steplike signals is confirmed by the detection results in the last phase for the ‘short latency’

⁷ For a group of crashes significantly closer than $\tau_{SC,det}$, only the largest of them can be detected in a single time trace. However, in another time trace, another of these edges can be larger. Thus, detection of the entire group of close-by crashes is still possible.

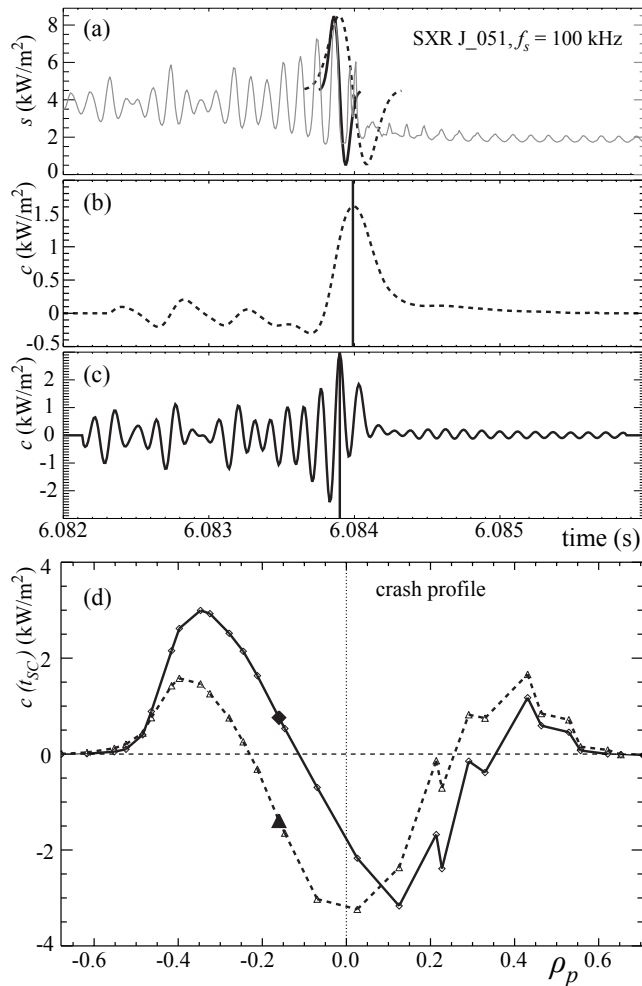


FIG. 10: Detection of a single sawtooth crash with the same $\tau_{SC,det} = 1$ ms but different kernel widths $\sigma_{k,1} = 0.095$ ms and $\sigma_{k,2} = 0.041$ ms, corresponding to resonance frequencies $f_{res,1} = 1.67$ kHz and $f_{res,2} = 3.9$ kHz. (a) Time trace of SXR channel with the two kernels plotted at the detected edge time point, (b) convoluted signals for $\sigma_{k,1}$ and (c) for $\sigma_{k,2}$, (d) crash profiles for both kernels. The position of the SXR channel shown in (a) is marked by filled symbols.

settings (dots in bottom part of figure 8). The other two methods do not even detect these large sawtooth crashes reliably.

Sawtooth detection at $\|l_f\| \approx 14$ is a special case because the results depend strongly on the signal type (steplike or ramplike) for the POSR method, but hardly for the kurtosis method. For the same time trace with sawteeth, the signal type for the detector depends on the selected $\tau_{SC,det}$: A sawtooth period of 11 ms will give fully ramplike signals for $\tau_{SC,det} = 11$ ms but nearly steplike signals for $\tau_{SC,det} = 4$ ms, because the latter analyses only a small fraction of the sawtooth cycle. In figure 11 we therefore compare the detection results for the POSR and the kurtosis method

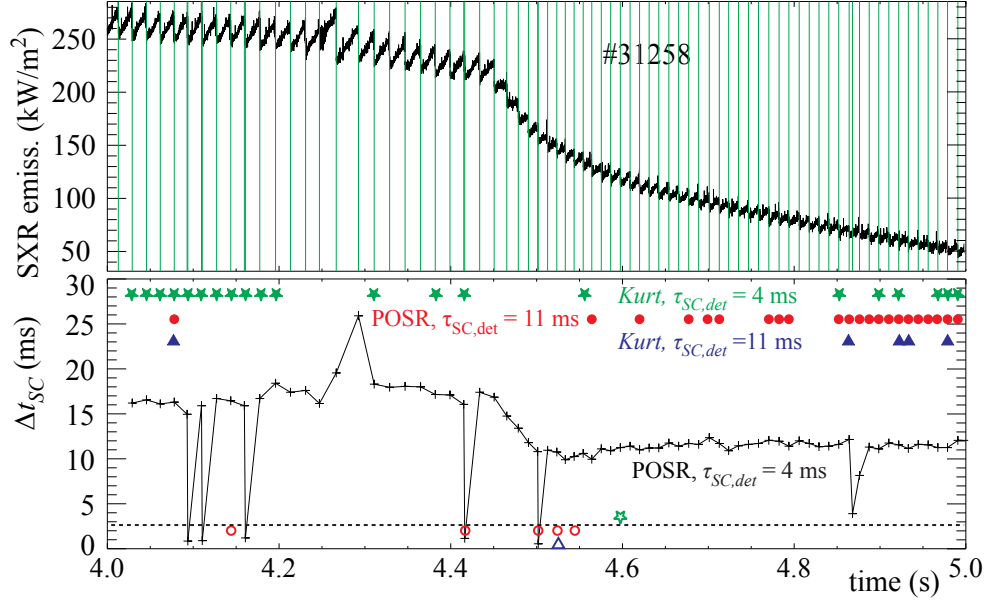


FIG. 11: Sawtooth detection at $\|l_f\| = 14$ ($f_s = 20$ kHz). Top: SXR time trace of discharge #31258 with detection results for POSR method with ‘11 ms steplike’ setting (which has $\tau_{SC,det} = 4$ ms), bottom: distances to previous detected crash. For the other three detection results - ‘11 ms ramplike’ setting for the POSR method and both settings for the kurtosis method - only the missed detections (i.e. false negatives) are marked by filled symbols and the false positives by open symbols.

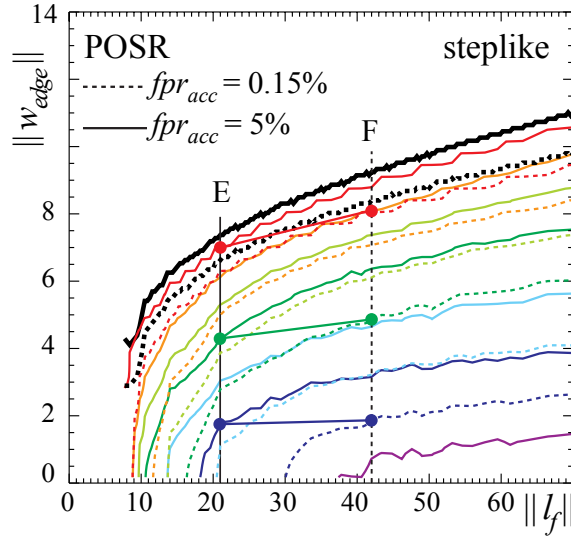


FIG. 12: Central graph of figure 6 with $\|l_f\| = 21$ and 42 marked as solid and dashed vertical lines, respectively. For some selected edge heights, the detection limits relevant for trace E ($\|l_f\| = 21$, $fpr_{acc} = 5\%$) and trace F ($\|l_f\| = 42$, $fpr_{acc} = 0.15\%$) in figure 8 are compared (connected large dots).

for different $\tau_{SC,det}$, and thus signal type, for another discharge (#31258, with pure NBI heating, switched off at 4.4 s). For $\tau_{SC,det} = 4$ ms, all crashes are nearly steplike, while for $\tau_{SC,det} = 11$ ms the second half of the crashes is fully ramplike. As expected from figure 6 (middle row, left), detection with the POSR method for steplike crashes is very good, while many ramplike crashes are missed at $\|l_f\| = 14$ (red dots in figure 11). The kurtosis method performs much better for $\tau_{SC,det} = 11$ ms, i.e. ramplike signals, revealing that the failure of POSR is not only in identifying peaks that are closer than $\tau_{SC,det}$. The kurtosis method has similar detection performance for steplike and ramplike signals in the simulation (see figure 6 (bottom, left)), such that for ramplike signals, the kurtosis method can still detect sawteeth at $\|l_f\| = 14$, where POSR fails. However, the normalized edge width must be sufficiently small. In the first half of figure 11, many crashes have a duration in the order of 1 ms. This results in $\|w_{edge}\| \approx 1$ for $\tau_{SC,det} = 11$ ms, which is in the detectable range in the simulation, but in $\|w_{edge}\| \approx 2.5$ for $\tau_{SC,det} = 4$ ms, which does not allow reliable sawtooth detection for the simulated signals. These observations are thus directly predicted from the simulations.

Traces E and F in figure 8 represent the detection results for an acceptable false positive rate of $fpr_{acc} = 5\%$ with the ‘intermediate’ setting ($\|l_f\| = 21$) and for $fpr_{acc} = 0.15\%$ with the ‘small crashes’ setting ($\|l_f\| = 42$), respectively. Both detections use the POSR method with the same kernel width ($\sigma_k = 0.25$ ms, corresponding to $f_{res} = 0.65$ kHz). According to figure 12, the ‘small crashes’ setting with $fpr_{acc} = 0.15\%$ should perform similarly as, or better than, the ‘intermediate’ setting with $fpr_{acc} = 5\%$. However, in the ECRH phase the detection results for POSR are much better for larger fpr_{acc} (trace E) than for larger $\|l_f\|$ (trace F). The other two methods show similar behaviour, which confirms that an increase of $\|l_f\|$ in case of non-stationary signal evolutions is less beneficial for detection performance than in the simulation.

6. SUMMARY AND OUTLOOK

Sawtooth crashes in tokamaks with metal walls can show strong variations from one crash to the next in terms of inversion radius, amplitude and even in the direction of the central signal change. In ASDEX Upgrade, this irregular behaviour is connected to central wave heating, especially during ECRH. Soft X-ray (SXR) signals are more affected than ECE. However, large fluctuations of inversion radius and crash amplitude have been observed also for ECE measurements. In such scenarios, detection methods that rely on predefined or slowly changing crash characteristics are not suited to reliably detect sawtooth crashes.

We designed an automatic sawtooth crash detection method with self-adjusting thresholds, unprejudiced in terms of sawtooth evolution and able to detect also very small crashes under a wide range of plasma conditions. This is done in three steps: The first step is an edge detection in individual time traces by convolution with a filter being shaped as the Gaussian first derivative (equation (1)), which results in a peak at the crash time point. The convoluted signal is evaluated in a calculation frame around the peak. For a fully adaptive edge detection, we use peak thresholds based on the noise level in the vicinity of the crash.

A twofold combination of detections over a range of diagnostic channels follows. First, a detection must prove its significance by appearing in at least two traces within a small time window. Then the spatial distribution is evaluated (see equations (10) to (12)) in terms of generic crash characteristics, such as relative amplitude with respect to the precrash values, symmetry with respect to the magnetic axis and a significant redistribution (i.e., both, gains and losses are observed). For the last step, a signal set which covers the mixing region well, is required. The first two steps can be performed with a smaller set of signals. In order to be independent of the ECE cut-off electron density restriction, we based the detection on SXR signals, as they are available in basically all scenarios. However, the method should be adaptable to other diagnostic methods, like e.g. ECE, provided the temporal resolution and spatial coverage are sufficient.

The effect of edge detection parameter choice has been simulated with synthetic signals. Three methods for evaluation of the convoluted signals have been compared (equations (6) to (9)). A method that corrects the noise reference for the influence of the peak (POSR) shows the best results, followed by the method that uses the kurtosis of the convoluted signal. The latter is a measure for the significance of the peak in its noisy background. The simulation considers the crash height and its duration but only with white noise. The presence of a sawtooth precursor thus necessitates consideration of specific perturbation frequencies in addition to the simulation results.

Depending on boundary conditions and preferences, suitable detection parameter settings can be extracted. Detection of very small sawtooth crashes requires a sufficiently large ratio of the frame length to the width of the convolution kernel. Extended crash duration or low perturbation frequencies require a large kernel width. The sum of frame length and kernel length (which is proportional to the kernel width), is limited by the minimum distance between sawtooth crashes and by the desired latency for real-time applications. Detection of very small crashes or crashes with a low-frequency precursor will conflict with very short latencies. Thus, for future real-time applications it must be decided whether the detection should be very fast and possibly miss some crashes, or if it is allowed to take few milliseconds longer but detect reliably virtually each sawtooth

crash. The same applies for the choice of acceptable false positive rate, fpr_{acc} . Few false positives after edge detection are preferable, since the profile analysis has to be performed for all detected edge time points. However, we have shown that the detection in a phase with inverted crashes can suffer from a low fpr_{acc} . To overcome these restrictions, different settings could be used in parallel. The parameter choices listed in table I were applied to ASDEX Upgrade SXR data with strongly different sawtooth features. The results confirm to a large extent our expectations based on the simulation and frequency considerations.

In the future, we will analyse differences in detection for the two more successful methods (POSR and kurtosis) and for different parameter sets on a broad basis of discharges. Application to two or more SXR cameras simultaneously could reduce sawtooth crash rejections. This will use the fact that the asymmetry of the crash profile due to a slow precursor with odd symmetry depends on the viewing angle. A real-time implementation of the algorithm with the shortly available real-time ASDEX Upgrade SXR data will be started within 2017. The presented automated sawtooth detection algorithm already routinely provides sawtooth crash times which are used to support equilibrium reconstruction [17].

Acknowledgments

The authors thank L. Giannone and H. Zohm for their valuable contributions. This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

-
- [1] R. Hastie, *Astrophysics and Space Science* **256**, 177 (1997).
 - [2] T. C. Hender, J. C. Wesley, J. Bialek, A. Bondeson, A. H. Boozer, R. J. Buttery, A. Garofalo, T. P. Goodman, R. S. Granetz, Y. Gribov, et al., *Nuclear Fusion* **47**, 128 (2007).
 - [3] O. Sauter, E. Westerhof, M. L. Mayoral, B. Alper, P. A. Belo, R. J. Buttery, A. Gondhalekar, T. Hellsten, T. C. Hender, T. Howell, D F ans Johnson, et al., *Physical Review Letters* **88**, 105001 (2002).
 - [4] F. Porcelli, D. Boucher, and M. N. Rosenbluth, *Plasma Physics and Controlled Fusion* **38**, 2163 (1996).
 - [5] M. van Berkel, G. Witvoet, M. de Baar, P. Nuij, H. ter Morsche, and M. Steinbruch, *Fusion Engineering and Design* **86**, 2908 (2011).
 - [6] H. van den Brand, M. R. de Baar, M. van Berkel, T. C. Blanken, F. Felici, E. Westerhof, M. Willensdorfer, ASDEX Upgrade Team, and EUROfusion MST1 Team, *Plasma Physics and Controlled Fusion* **58**, 075002 (2016).

- [7] M. Lennholm, L.-G. Eriksson, F. Turco, F. Bouquey, C. Darbos, R. Dumont, G. Giruzzi, M. Jung, R. Lambert, R. Magne, et al., *Fusion Science and Technology* **55** (2009).
- [8] M. Lennholm, T. Blackman, I. T. Chapman, L. G. Eriksson, J. P. Graves, D. F. Howell, M. de Baar, G. Calabro, R. Dumont, M. Graham, et al., *Nuclear Fusion* **51**, 073032 (2011).
- [9] T. P. Goodman, F. Felici, O. Sauter, J. P. Graves, and the TCV team, *Physical Review Letters* **106**, 245002 (2011).
- [10] I. G. J. Classen, J. E. Boom, W. Suttrop, E. Schmid, B. Tobias, C. W. Domier, J. N. C. Luhmann, A. J. H. Donne, R. J. E. Jaspers, P. C. de Vries, et al., *Review of Scientific Instruments* **81**, 10D929 (2010).
- [11] A. Gude, M. Maraschek, C. Angioni, J. Stober, and ASDEX Upgrade Team, in *Europhysics Conference Abstracts (CD-ROM, Proc. of the 37th EPS Conference on Plasma Physics, Dublin, Ireland, 2010)* (2010), vol. 34A of *ECA*, p. P4.124, URL <http://ocs.ciemat.es/EPS2010PAP/pdf/P4.124.pdf>.
- [12] M. Sertoli, J. M. Garcia-Regana, T. Odstrcil, C. Angioni, H. Smith, Y. Turkin, W. A. Cooper, A. de Bustos, D. Vezinet, ASDEX Upgrade team, et al., in *Europhysics Conference Abstracts (CD-ROM, Proc. of the 42th EPS Conference on Plasma Physics, Lisbon, Portugal, 2015)* (2015), vol. 39E of *ECA*, p. O4.129, URL <http://ocs.ciemat.es/EPS2015PAP/pdf/O4.129.pdf>.
- [13] A. Gude, M. Maraschek, P. Eulenberg, V. Igochine, and ASDEX Upgrade Team, in *Europhysics Conference Abstracts (CD-ROM, Proc. of the 42th EPS Conference on Plasma Physics, Lisbon, Portugal, 2015)* (2015), vol. 39E of *ECA*, p. P1.122, URL <http://ocs.ciemat.es/EPS2015PAP/pdf/P1.122.pdf>.
- [14] Y. F. Baranov, C. D. Challis, J. Ongena, B. Alper, G. Arnoux, P. Buratti, T. Gerbaud, D. Keeling, V. Kiptily, J. Mailloux, et al., *Nuclear Fusion* **52**, 023018 (2012).
- [15] J. A. Canny, *IEEE Trans. Pattern Analysis and Machine Intelligence* **8**, 679 (1986).
- [16] Z. Pietrzyk et al., *Nuclear Fusion* **39** (1999).
- [17] R. Fischer, in *Europhysics Conference Abstracts (CD-ROM, Proc. of the 43th EPS Conference on Plasma Physics, Leuven, Belgium, 2016)* (2016), vol. 40A of *ECA*, p. P1.016, URL <http://ocs.ciemat.es/EPS2016PAP/pdf/P1.016.pdf>.