Sawtooth pacing by real-time auxiliary power control in a tokamak plasma

T P Goodman, F Felici, O Sauter, J P Graves and the TCV Team

Centre de Recherches en Physiques des Plasmas,

Ecole Polytechnique Fédérale de Lausanne (CRPP - EPFL), Lausanne, Switzerland, CH-1015

(Dated: June 6, 2011)

In the standard scenario of tokamak plasma operation, sawtooth crashes are the main perturbations that can trigger performance-degrading, and potentially disruption-generating, neoclassical tearing modes. This Letter demonstrates sawtooth pacing by real-time control of the auxiliary power. It is shown that the sawtooth crash takes place in a reproducible manner shortly after the removal of that power, and this can be used to precisely prescribe, i.e., pace, the individual sawteeth. In combination with preemptive stabilization of the neoclassical tearing modes, sawtooth pacing provides a new sawtooth control paradigm for improved performance in burning plasmas.

In high performance tokamaks, the plasma β (the ratio of plasma pressure to magnetic pressure) is often limited by metastable magnetohydrodynamic (MHD) instabilities known as neoclassical tearing modes (NTMs) [1]. These modes are of special concern for fusion-reactor-grade tokamak plasmas since the islands generated in the confining magnetic topology grow to a size, even at relatively low β, that they must be mitigated, or avoided altogether, to allow sufficiently economic power generation. Fortuitously, these modes will self-stabilize if the island is reduced below a certain size, and moreover they require a seed island to begin to grow. The main perturbation that triggers these modes in the standard tokamak scenario is the sawtooth core relaxation event [2, 3] that can occur when the safety factor q is below one (q is the rate of change of the toroidal flux with the poloidal flux). At the "crash" of the sawtooth, particle and magnetic energy are redistributed from inside to outside the surface where $q = 1$ and the liberation of this energy can seed MHD modes at resonant surfaces where $q = m/n$ is rational (m and n are the integer poloidal and toroidal mode numbers, respectively), in particular at the $q = 3/2$ and $q = 2/1$ surfaces. Large 2/1 modes can also lead to plasma current disruptions, which might damage a reactor.

NTM avoidance is concerned with eliminating or reducing the triggering mechanisms. The scenario foreseen in ITER to reach the optimal fusion power, the so-called standard scenario, is a sawtoothing plasma with long sawtooth periods resulting from the strong fast-partical stabilization inherent in burning plasmas [4, 5]. It has been shown in many tokamaks that the crashes of long-period sawteeth can trigger NTMs even at low β [6, 7]. Therefore, the control of the sawtooth period is crucial for NTM avoidance and emphasis has been placed on shortening the period. It has been demonstrated in the tokamak à configuration variable (TCV) [8] and Tore Supra [9] that the sawtooth period can be controlled by feedback-positioning very-locally absorbed EC waves relative to the $q = 1$ surface. However, the dependence of the sawtooth period on the the absorption location is extremely non-linear. We propose a new approach for the control of the sawtooth period, that of sawtooth pacing by power control. The method provides a more robust control and opens new possibilities since, as demonstrated in this letter, it can regulate when the next sawtooth crash will occur very accurately, without relying on real-time simulations as was suggested in [10].

In contrast to NTM avoidance, as advocated in this letter, mitigation of NTMs relies on removing or "healing" the islands, either as they grow or, if growth is too rapid, once they are saturated [10] (though in future large tokamaks, modes may lead to a disruption before they saturate). Stabilization relies primarily on providing current in the island by electron cyclotron heating (ECH) or current drive (ECCD) [11]. ECCD is the ideal tool for stabilizing NTMs since the current can be localized at the resonant surface where the island exists, and made more efficient by modulation in phase with the O-point of the (rotating) island [13] as the island becomes small. For this reason, a primary function of the EC upper launchers in the ITER tokamak is to stabilize the $q = 2/1$ and $q = 3/2$ NTMs [12]. Applying *preemptive* ECCD continuously at the rational surface where the mode is *expected* to appear has been shown to avoid the growth of NTMs while using less power than that required to shrink the mode once it has grown [14]. Sawtooth pacing that successfully predicts the time of the upcoming sawtooth crash permits preemptive ECCD to be activated prior to the crash and sustained only until any potential modes generated by the crash have been avoided. Thus, continuous ECCD might not be required, as long as NTMs do not (as expected in the ITER standard scenario [3]) appear between sawtooth crashes.

The experiments presented in this letter were made possible by the recently commissioned TCV digital real-time control (RTC) system [15]. One node of this control system acquires 64 chords of a Dual Multiwire Proportional X ray detector (DMPX), measuring line-integrated soft-X ray intensity as seen from the machine floor. A
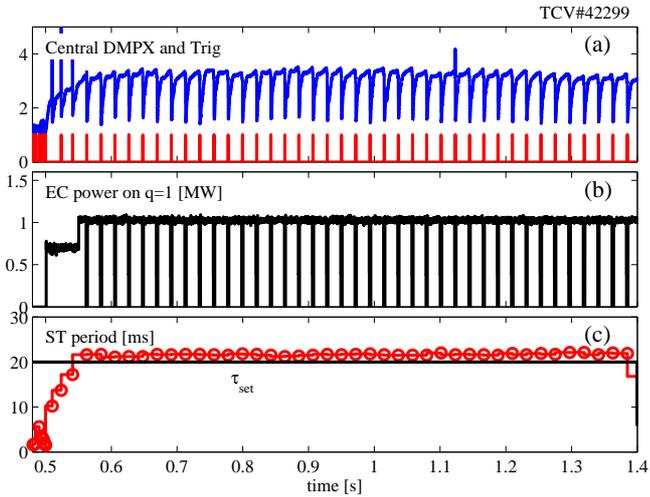
FIG. 1. (a) The x-ray measurements (DMPX) are used to detect the sawtooth crash and initiate the stabilizing ECCD from 0.55 s to the end of the discharge. (b) The EC power is removed after 20 ms and restarted at the next sawtooth crash throughout the discharge. (c) The sawtooth period is thereby paced at 21.6 ± 0.2 ms. The RTC resolution is 0.1ms.
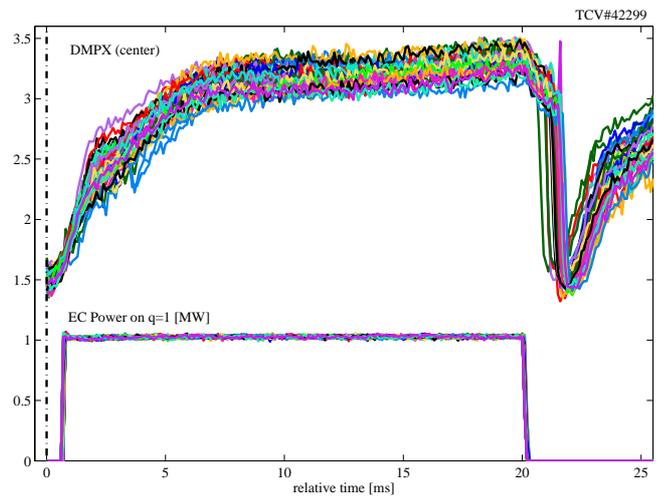

FIG. 2. The central x-ray measurements (DMPX) from the 38 consecutive sawteeth are overlaid showing the high level of reproducibility of the crashes, the droop in temperature at the power removal and that the sawteeth are much longer than the confinement time. The associated ECCD pulse waveforms ($\tau_{set} = 20$ms) are shown below with relative time measured from the detection of the (preceding) sawtooth crash, indicated by the dashed line.

real-time control algorithm runs on an embedded X86 Linux PC with a cycle time of 0.1 ms. The seven central chords of the DMPX are pre-filtered and averaged, yielding a signal with the traditional "sawtooth" shape. This signal is then passed to a sawtooth crash detector algorithm that detects the occurrence of a sawtooth crash, from one cycle to the next, by checking criteria based on the current sample value and its difference from the previous sample. The trigger signal generated by each sawtooth crash is used to reset a running timer. When the timer exceeds a pre-set time τ_{set}, the EC power is reduced or switched off. When the next sawtooth crash is detected, the EC power is reinstated and the cycle begins anew.

The optimal location for stabilization in a limited L-mode plasma with fixed EC mirror position and plasma shape (elongation = 1.52, triangularity = +0.33) was determined by performing a scan of the toroidal field, B, together with the plasma current I_p. The same mirror position and shape are kept for the experiments reported herein but the absorption location is held constant, with $B = 1.19$ T and $I_p = 0.33$ MA giving $q_{edge} = 2.7$. Figure 1 shows the RTC pacing of the sawtooth period, τ_s, through controlled removal of the stabilizing ECCD near the $q = 1$ surface. The crash detection signal triggers stabilizing ECCD pulses of a chosen length, at each crash, once the RTC is engaged (at 0.55 s), as described above. In this experiment $\tau_{set} = 20$ ms resulting in sawteeth that are regulated at $\tau_s = 21.6 \pm 0.2$ ms. The sawtooth crash was observed to occur at a reproducible interval after the removal of the power.

The central DMPX traces, and the controlled EC power traces, for each sawtooth cycle are overlaid in Fig. 2 (synchronized in time to the previous detected crash). The relatively flat x-ray signal evolution indicates that τ_s is much longer than the electron energy confinement time, τ_{E_e}. The x-ray signal change is dominated by the electron temperature during the sawtooth cycle, which is seen to decrease at the time the power is removed. The variation in the sawtooth period is very small and the crash occurs shortly after the power removal indicating that the sawteeth are close to marginal stability. (Note that the gyrotron voltage ramps up for 0.7 ms and ramps down over 0.3 ms, leading to the delay seen in the start of the power traces in Fig. 2.)

This sawtooth control method has been used to avoid disruptions caused by modes generated at the crash of a long-period sawtooth by choosing $\tau_{set} = 20$ ms, which is well below the period that typically triggers a disruption (τ_s somewhat greater than 30 ms in these plasmas). When the sawtooth period is extended by ramping τ_{set}, the response to the change in the ECCD pulse length occurs immediately: each subsequent crash occurs at the new τ_{set} plus the delay of ~ 2 ms, which remains nearly constant in this experiment. Figure 3 shows the increase in τ_{set}, and thus τ_s, starting at 1.0 s. At the shorter sawtooth periods, post-crash MHD activity is seen on the Mirnov coils at each crash (Fig. 3c). Near $t = 1.13$ s, a tearing mode is triggered which self-stabilizes before the next large sawtooth crash via the formation and interplay of several modes; smaller crash-like events are seen in the DMPX signal (Fig. 3a), and influence the evolu-
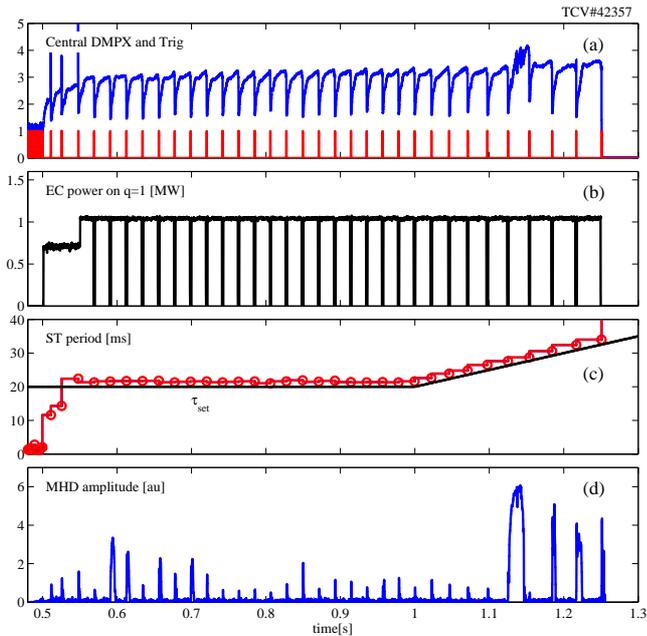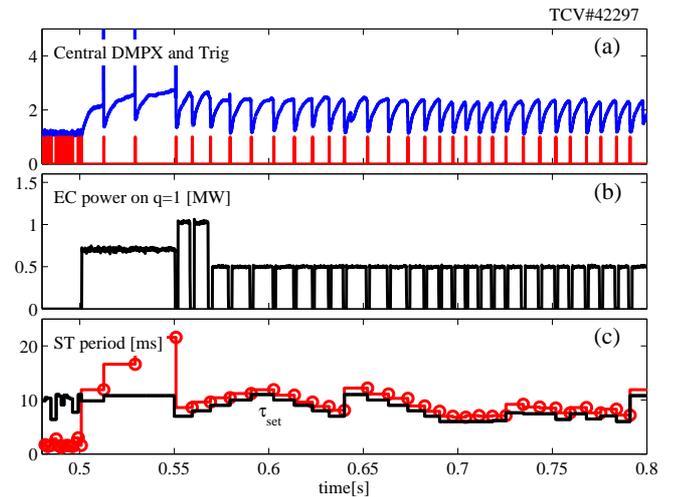
FIG. 4. Sawtooth crashes are controlled on an individual basis by the RTC from 0.55 s onward: (a) DMPX signal and sawtooth detector output, (b) the stabilizing ECCD power and (c) the requested τ_{set} and resulting τ_s.

FIG. 3. (a) DMPX signal and sawtooth detector output, (b) the resulting ECCD power pulses, (c) τ_{set} and τ_s, and (d) the rms amplitude of a Mirnov coil. The discharge ends in a disruption when a sufficiently long sawtooth period is reached.

tion of the modes present. The triggering of the modes by the long sawtooth crashes and their interplay will be the subject of a separate publication; our concern here is the control of the sawtooth period itself which governs the appearance of the modes. Finally, at $t = 1.25$ s, a sawtooth crash occurs with $\tau_s = 34$ ms, a large 2/1 mode is triggered and the plasma disrupts.

Figure 4 establishes that the sawtooth period of *individual* sawteeth can be controlled. At each sawtooth crash, τ_{set} is modified to generate a "staircase-like" series of periods - both increasing and decreasing by small steps, as well as larger jumps, from one sawtooth crash to the next - proving that extremely fine control of the period is possible with this methodology. Furthermore, it provides an experimental demonstration that, due to the reconnection at each crash, a new initial condition is set that is effectively independent of the pulse history. The RTC begins at 0.55 s and the long sawtooth periods prior to this time approach the "natural" period of the fully stabilized sawteeth, $\tau_{stabilized}$, with an ECCD power level of 0.75 MW (Fig. 4b). This highlights the fact that this type of control is only possible for $\tau_{set} < \tau_{stabilized}$, where $\tau_{stabilized}$ depends on the chosen auxiliary heating source and its power level. It is important to note that the first two power pulses are at 1.0 MW full power, as in Fig. 3b, and that the attainable sawtooth period at this power is, therefore, at least 34 ms. Nevertheless, removal of the stabilizing power reduces the sawtooth period by nearly a factor of 4, to 9 ms (at the smallest

τ_{set} used in the experiment). An even larger reduction is likely if τ_{set} is reduced further. At 0.57 s, one of the EC power sources is lost decreasing the total input power by a factor of two. The sawtooth period control continues to function properly even given this strong perturbation since the requested sawtooth period is still smaller than $\tau_{stabilized}$ for one power source. If the power level had been too low to extend the sawtooth period beyond τ_{set}, the algorithm would simply have maintained the power on and the sawtooth crashes would have occurred at the uncontrolled, "natural" rate for that power level. Despite the loss of half of the EC power the sawtooth periods are still individually controlled in our case.

We have also seen that when the power is partially (as opposed to completely) removed after τ_{set}, the next crash occurs after a longer delay and with lower reproducibility (though still with only a small jitter). Thus, with $\tau_{set} = 7$ ms and the full power (0.45 MW) completely removed, $\tau_s = 8.0 \pm 0.2$ ms, whereas when the power is reduced to 0.2 MW between full power plateaus $\tau_s = 9.4 \pm 0.8$ ms. This experiment simulates the intentional control of only some of several, combined, stabilizing sources, or their unintentional loss (e.g. due to malfunction). Finally, when the modulation frequency and duty cycle are kept constant at values close to those of the two above-mentioned experiments (8 ms of full power followed by 2ms off) it is observed that the sawtooth period was not controlled and τ_s fluctuated by more than a factor of two; that is, the sawteeth are not simply "locking" to the modulation frequency.

In ITER, as shown in [16], it is expected that the sawtooth period can be increased by a factor of 1.5 (compared to the period resulting from fast-ion stabilization)

using 13.3 MW of ECCD when the absorption location is feedback controlled to be at an optimum location near the $q = 1$ surface using the EC upper launchers [12]. In the ITER simulations, a classical model [5, 17] for the sawtooth crash criteria and post-crash profiles was incorporated in a transport code to determine the stabilizing potential of ECCD. Despite being heuristic, this model is of significant practical importance (see, for example, the discussion in Ref. [7]): previous calculations, based on the model, successfully simulated detailed sawtooth stabilization experiments using ECCD in the tokamak à configuration variable (TCV) [18] and neutral beam injection (NBI) [19] and ion cyclotron current drive (ICCD) [20] in the joint european torus (JET). The many criteria of the model depend critically on the magnetic shear $[s = (r/q)dq/dr]$ at the $q = 1$ surface, s_1, and changes of sawtooth period with deposition location (and/or current drive direction) result, primarily, from changes in the temporal derivative of s_1.

An important difference in ITER, compared with TCV, is the presence of large populations of highly energetic ions, comprising fusion alpha particles and those generated by NBI and ion cyclotron resonance heating (ICRH) systems. It is known that fast ion populations can greatly lengthen the sawtooth period. The magnetically trapped fractions of centrally-peaked fast-ion distributions increase the threshold for instability, and within the framework of the stability criteria of Ref. [5], a sawtooth crash is delayed until the shear has time to reach the higher threshold. The instability threshold can also be raised by energetic passing ions when centrally deposited neutral beams are injected in the co-current direction [21], or with co-current propagating ion cyclotron waves resonating with a low concentration minority ion population close to the $q = 1$ surface [4, 6]. Recent simulations show that the effects of the planned NBI and ICRH heating systems on the internal kink mode stability in ITER will be comparable to the effect of fusion alpha particles. It can therefore be expected that removal of the stabilizing NBI or ICRH power would result in a sawtooth crash after a relatively short delay, comparable with the slowing down of the fast ions, as was demonstrated in early JET ICRH experiments [22]. For this reason, we expect the sawtooth pacing paradigm put forward in this letter to apply not only to the localized ECCD actuator employed here, but also to NBI and/or ICRH heating methods.

Real-time sawtooth pacing addresses the main issue of the alternative approach – that of predicting the moment of the next sawtooth crash. It provides a purely experimentally-based real-time approach to determining when the next crash will occur, in contrast to the modelling-based real-time calculations referred to in Ref. [10] (which can nevertheless provide useful physics-based inputs to the control algorithms). Off-axis preemptive ECCD could then be applied before the sawtooth crash

occurs and removed some time after the accompanying seeding perturbation has passed. In this way, the time during which off-axis ECCD power is used may be minimized. Given the slow time scales of large burning plasma experiments such as ITER and the rapid movement of the EC actuators, it is conceivable that the same EC power sources and launchers might be used for both sawtooth stabilization and then preemptive NTM control at a different location; leaving the rest of the EC power available for central heating and current drive. In addition, sawtooth pacing using RTC switch-off of other actuators like ICRF or NBI should also be possible since they have also been employed to strongly modify the sawtooth period.

These experiments demonstrate that the moment of the sawtooth crash can be reliably and precisely controlled, changed and predicted from one sawtooth period to the next. The modes that are generated at the sawtooth crash, when the period is excessively long, can be avoided by choosing to induce the crash before that period is reached. Alternately, in the case of NTMs that are seeded by the crashes of stabilized sawteeth, knowledge of when the seeding event will occur can be used to apply preemptive avoidance, or active mitigation, techniques in an optimal way. Investigation of optimum preemptive NTM control is the subject of continued research at TCV.

* timothy.goodman@epfl.ch

[1] O. Sauter et al., Phys. Plasmas **4**, 1654 (1997).

[2] O. Sauter et al., Phys. Rev. Lett. **88**, 105001 (2002).

[3] T.C. Hender et al., Nucl. Fusion **47** S128 (2007).

[4] J. P. Graves et al., Phys. Rev. Lett. **102**, 065005 (2009).

[5] Porcelli F, Boucher D and Rosenbluth M Plasma Phys. Control. Fusion **38**, 2163 (1996).

[6] J. P. Graves et al., Nucl. Fusion **50**, 052002 (2010).

[7] I. T. Chapman, Plasma Phys. Control. Fusion **53**, 013001 (2011).

[8] J. I. Paley et al., Plasma Phys. Control. Fusion **51**, 055010, (2009).

[9] M. Lennholm et al., Phys. Rev. Lett. **102**, 115004 (2009).

[10] O. Sauter et al., Plasma Phys. Control. Fusion **52**, 025002 (2010).

[11] G. Gantenbein et al., Phys. Rev. Lett. **85**, 1242 (2000).

[12] M.A. Henderson et al., Nucl. Fusion **48**, 054013 (2008).

[13] M. Maraschek et al., Phys. Rev. Lett. **98**, 025005 (2007).

[14] K. Nagasaki, A. Isayama, S. Ide and JT-60 Team Nucl. Fusion **43**, L7 (2003).

[15] J.I. Paley et al., Real Time Conference (RT), 2010 17th IEEE-NPSS , p. 1 (IEEE, New York, 2010).

[16] C. Zucca, Ph.D. thesis, EPFL [Report No. 4360 Lausanne, 2009, http://library.epfl.ch/theses/?nr=4360].

[17] O.Sauter *et al.*, Theory of Fusion Plasmas, ISPP-18, ISBN 88-7794-167-7, p. 403 (1998).

[18] C. Angioni *et al.*, Nucl. Fusion **43**, 455 (2003).

[19] C. Angioni *et al.*, Plasma Phys. Control. Fusion **44**, 205 (2002).

[20] M. J. Mantsinen *et al.*, Plasma Phys. Controlled Fusion **44**, 1521 (2002).

[21] J. P. Graves, Phys. Rev. Lett. **92**, 185003 (2004)

[22] D. J. Campbell *et al.*, 15th EPS conf., Dubrovnik, 1988, ECA, Vol. 12B, p. 377.