2012 American Control Conference
Fairmont Queen Elizabeth, Montréal, Canada
June 27-June 29, 2012

Robust adaptive control of the sawtooth instability in nuclear fusion

J.J. Bolder[1], G. Witvoet[1,3,4], M.R. de Baar[1,3], N. van de Wouw[2],
M.A.M. Haring[2], E. Westerhof[3], N.J. Doelman[4] and M. Steinbuch[1]

Abstract— **The sawtooth instability is a repetitive phenomenon occurring in plasmas of tokamak nuclear fusion reactors. Experimental studies of these instabilities and the effect they have on the plasma (notably the drive of secondary instabilities and consequent performance reduction) for a wide variety of plasma conditions is an important line of study in nuclear fusion research. Variations in the plasma conditions have a significant influence on the dynamical behavior of the sawtooth instability. Therefore, this paper presents the design of a sawtooth period controller which is robust against such variations. The controller is from a class of adaptive controllers better known as extremum seekers. In this technique, a cost function in terms of the desired sawtooth period is optimized on-line. The Extremum Seeking Controller (ESC) is model-free and is therefore inherently robust against model uncertainty. Simulations show that the controller is robust against variations in plasma parameters, delay in the sawtooth crash detection, and noise on the in- and outputs of the process. Because of its robustness, ESC is a promising candidate strategy for a wide range of fusion-related control problems with high model uncertainty.**

I. Introduction

Nuclear fusion is a promising solution to the global energy problem [1], although it comes with a number of technical challenges. In Tokamaks [2], the plasma is confined by strong magnetic fields in a toroidal shaped vessel. The plasma is subject to various instabilities associated with a reorganization of the magnetic topology of the plasmas. The sawtooth instability [3], [4] is a prime example of these. The sawtooth instability is characterized by a gradual evolution of the plasma variables e.g. temperature. The evolution is reset during a fast, crash-like, event. In the core of the plasma, the temperature slowly increases, and at a crash the temperature drops on a very short timescale. Consequently the core temperature measurements have a sawtooth-like shape, hence the name, sawtooth instability.

The sawtooth instability periodically mixes the plasma core [2]–[4]. This could provide a natural mechanism to regulate the exhaust of Helium ash [5]. Also the refueling (the influx of fusion fuels deuterium and tritium) could be enhanced by the sawtooth. On the other hand, the sawtooth instability can trigger secondary instabilities, which reduce the operational performance of the fusion plasmas and could lead to disruption [6], [7]. Control of the sawtooth period, i.e., the interval between two subsequent crashes, is necessary to avoid triggering other instabilities while concurrently refreshing the plasma core.

The sawtooth period can be influenced by driving a non-inductive current in the vicinity of the instability. This can be done by coupling electron cyclotron waves into the plasma. Resonant absorbtion of these waves yields a localized modification of the electron distribution and current. Co-Current drive (parallel to the main current in the plasma) close to the plasma center leads to shortening of the sawtooth period, whereas injecting outward the center lengthens the sawtooth period. The location of the deposition can be varied using a mechanical steerable mirror, the launcher [8]–[10]. The launcher has its own motion control system which we will consider as part of the process. Usually, sawteeth are measured using soft x-rays or Electron-Cyclotron Emission. An algorithm to accurately determine the sawtooth period from such temperature measurements is discussed in [11].

Closed-loop sawtooth controllers based on linear control theory have been developed on TCV (Tokamak configuration variable, Lausanne, Switzerland) [12] and Tore Supra (Cadarache, France) [13], [14]. The successful tracking of a sawtooth period reference has been demonstrated. However, there is no analysis of the dynamic behaviour of the sawtooth period, which is essential to guarantee closed-loop stability and tracking performance. In [12], it is shown that the sawtooth period controller is stable for small period sawteeth, and is actually unstable for larger period sawteeth, where the controller introduces heavy oscillations in the launcher angle. This stability problem is addressed in [15], where the dynamics of a sawtooth model with constant plasma parameters is identified in a selected region of the operating range by specifying limits on the launcher angle. A stabilizing linear feedback controller is designed, and the tracking of sawtooth period references is demonstrated.

However, changes in plasma parameters significantly change the behavior of the sawtooth instability and may endanger stability of the controllers. A different sawtooth controller is proposed in [16], which is from a class of adaptive controllers, commonly known as an extremum seeker [17]. In [16], the sawtooth period on TCV is maximized using Electron Cyclotron Current Drive (ECCD). The proposed extremum seeking controller (ESC) can only maximize the

[1] Eindhoven University of Technology, Dept. of Mechanical Engineering, Control Systems Technology group, P.O. Box 513, 5600 MB Eindhoven, The Netherlands, j.j.bolder@tue.nl
[2] Eindhoven University of Technology, Dept. of Mechanical Engineering, Dynamics and control group
[3] FOM-DIFFER, Association EURATOM-FOM, Trilateral Euregio Cluster, P.O. Box 1207, 3430 BE Nieuwegein, The Netherlands
[4] TNO Technical Sciences, Opto-mechatronics, P.O. Box 155, 2600 AD Delft, The Netherlands
The work in this paper has been performed in the framework of the NWO-RFBR Center of Excellence (grant 047.018.002) on Fusion Physics and Technology. This work, supported by the NWO, ITER-NL and the European Communities under the contract of the Association EURATOM/FOM, was carried out within the framework of the European Fusion Programme. The views and opinions expressed herein do not necessarily reflect those of the European Commission.

sawtooth period, and the controller structure has not been optimized. It was also reported that the controller parameters were difficult to select. The ESC in [16] is based on [18], [19], where a specific controller structure is proposed. Recent developments on the subject of ESC [17], [20], [21] extend the work in [18], [19], and allow for the systematic design and tuning of an ESC.

Clear requirements on the sawtooth controller have not been formulated in the fusion community. Therefore, a sequence of closed and open loop sawtooth period controllers has been derived with varying performance and robustness [15], [22]–[24]. In this paper, the systematic design of an ESC is presented, tailored to the sawtooth control problem; based on [21], [25]. The controller can track any feasible time-varying sawtooth period reference, unlike [16], which broadens its practical ability. Moreover, the controller is shown to be robust against varying plasma parameters and disturbances such as noise on the current drive and sawtooth period, and delay in the sawtooth crash detection.

The model proposed in [15] is used as a case study to test the controller. Although we only benchmark the controller in simulation, the aim of this paper is to develop a control strategy that is readily applicable in practice. This implies that the controller must be robust against variations in the sawtooth behaviour and disturbances in order to function on different tokamaks.

This paper is organized as follows: in Section II the sawtooth control problem is discussed. Section III elaborates on the controller design; the tuning of the controller parameters is briefly addressed in Section IV. Simulation results are presented in Section V and finally the conclusion and future work is featured in Section VI. An extended version of this paper is found in [26].

II. Control problem formulation

In this paper, the sawtooth model as elaborately discussed in [15], is used to benchmark the presented control strategy. The variable to be controlled is the interval between two sawtooth crashes, the sawtooth period. There is no measurement of the period in between crashes, it is a discrete variable. Since the period is also the variable to be controlled, the controller is designed in a discrete-time fashion where the controller action is triggered by a crash; this concurs with the approach taken in [15].

The sawtooth period is influenced by injecting an electron cyclotron waves via the launcher. A schematic representation of the actuation principle, known as Electron Cyclotron Current Drive (ECCD), is depicted in Fig. 1. The accessible deposition locations are indicated with the dotted line. The deposition is due to a resonant process, and hence the current is driven locally. The current is driven in toroidal direction. There are two inputs to the process. The first input is the electron-cyclotron driven current $I_{\rm CD}$, which is kept constant. The second input is the actual launcher angle $\vartheta_{\rm a}$; the desired launcher angle is denoted by ϑ. The output of the sawtooth process is a set of measurements χ, for instance, soft x-rays or Electron Cyclotron Emission (ECE), from which a sawtooth period $\tau_{\rm s}$ is determined by a period detection algorithm [11].

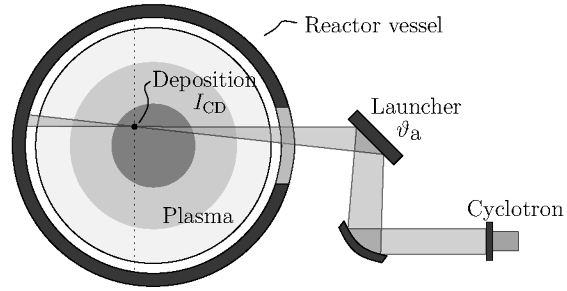

Fig. 1. Schematic representation of the Electron Cyclotron Current Drive (ECCD) system, showing a poloidal-radial cross-section of the toroidal reactor vessel. The current deposition is $I_{\rm CD}$, the deposition location is determined by the launcher angle ϑ.

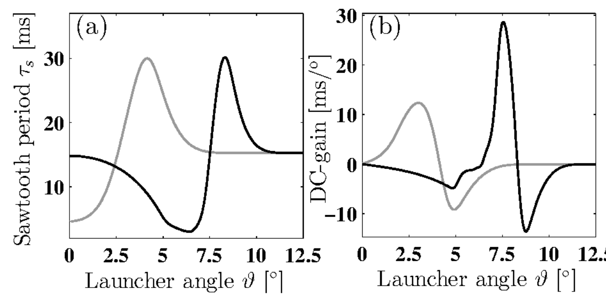

Fig. 2. Steady-state I/O map 1 (black) and 2 (gray) for a realistic change in plasma parameters (a), and the corresponding DC-gains (b).

The algorithm detects a sawtooth crash and triggers control action update.

As we will later argue, our proposed control strategy operates on a timescale that is slower than the timescale of the sawtooth dynamics. This implies that the desired launcher adjustments are so slow that the sawtooth period is always close to its steady-state value. We consider a steady-state input-to-output (I/O) map being the relation between the launcher angle ϑ and the sawtooth period $\tau_{\rm s}$ in steady-state. For controller design, it is relevant how the steady-state sawtooth period is affected when a small change in the desired launcher angle is applied. This is described by the DC-gain, which is the derivative of the I/O map with respect to the launcher angle. The sign of the DC-gain determines the control direction.

The I/O map and the DC-gain have been identified with the model for two realistic sets of plasma parameters, and are shown in Fig. 2. The variation in process behaviour is difficult to handle from a control point of view since the control direction actually changes, as plasma parameters change; this is indicated by the DC-gain in Fig. 2. Most classical control strategies as used in [12]–[15] rely on accurate knowledge of the DC-gain and process dynamics to ensure stability of the control loop. Herein, the sign of the DC-gain is not allowed to change; hence all these controllers are limited to a particular selected region of the operating range where the sign of the DC-gain is constant. Furthermore, plasma parameters often tend to drift during experiments, possibly destabilizing the control loop if the change in DC-gain is large enough.

We consider the following control problem. Design a discrete-time controller that, based on measurements of the sawtooth period, determines the desired launcher angle such that a desired sawtooth period is achieved, and is highly robust

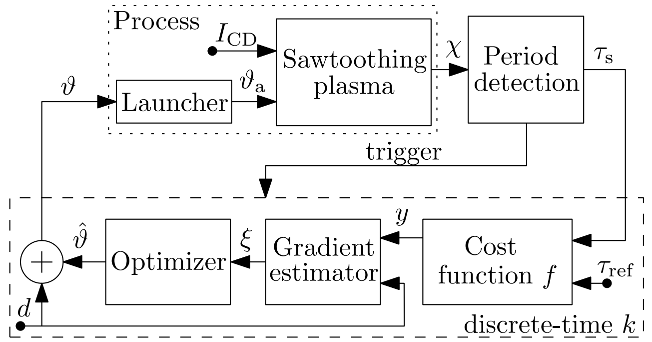

Fig. 3. Extremum seeking controller topology. The process is indicated with the dotted box and operates in continuous time, the controller is indicated with the dashed box and operates in discrete-time, controller action is triggered if a sawtooth crash is detected.

against variations in plasma parameters and disturbances.

A control strategy which is able to handle such changes is Extremum Seeking Control (ESC), since in this framework the DC-gain can be identified on-line. The next section discusses the design of such an ESC for the sawtooth control problem.

III. Extremum seeking controller design

Extremum seeking control is an adaptive control strategy that uses on-line optimization techniques to slowly drive a process to a desired operating point. The objective of the extremum seeker is finding the minimizer of a cost function f. For the sawtooth control problem this is finding the desired launcher angle ϑ that minimizes f such that a reference sawtooth period is tracked.

The block scheme in Fig. 3 shows the topology of the extremum seeking controller which is adapted to the sawtooth control problem. The discrete-time extremum seeking controller is indicated with the dashed box. The variable k is the crash counter and is the measure of discrete time.

The ESC consists of three subsystems: a cost function, a gradient estimator and an optimizer. The cost function is selected such that its function value y is minimal if the measured sawtooth period τ_{s} is equal to the reference sawtooth period τ_{ref}. The gradient estimator uses a perturbation signal d to estimate the gradients of the cost function with respect to $\hat{\vartheta}$, i.e., the estimate of the optimal launcher angle. The gradient estimate is denoted by ξ. The optimizer uses the gradient information to drive $\hat{\vartheta}$ to the minimizer of the cost function ϑ^*, which is typically unknown since the I/O-map is not known exactly. ESC requires that the I/O-map is stable [27], and that each input ϑ yields a unique output τ_{s}; the I/O map in Fig. 2, of which each point is indeed stable [15], indicates that the considered sawtooth model meets this criterion.

For this control strategy to function properly, it is essential to maintain a separation of the time-scales that each subsystem operates in [21]. The cost function is 'infinitely' fast, since it is static, so could be viewed as part of the process. The gradient estimator makes an on-line estimation of the gradients of the cost function with respect to the launcher angle. This estimation can only be performed correctly if the dynamics of the sawtooth instability have converged sufficiently close to their steady-state. The perturbation d used by the gradient estimator has to vary on a slower timescale than the process dynamics, such that the process is always operating close to its steady-state. Some gradient estimators need settling time due to internal filtering, others rely on a slow optimizer in order to work; hence the optimizer operates at the slowest timescale [17], [21]. The following subsections each discuss the design of the different subsystems.

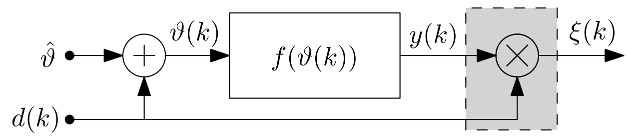

Fig. 4. Topology of a minimal gradient estimator, indicated in gray.

A. The cost function

The task of the extremum seeker is to find a launcher angle that minimizes a static cost function f. This cost function should be selected such that the system is in some desired operating condition when the function value is minimal. For the sawtooth control problem, the specification on the operating condition is a desired reference sawtooth period τ_{ref}. Therefore we propose the following cost function:

$$f(\tau_{\mathrm{s}}(k), \tau_{\mathrm{ref}}) = (\tau_{\mathrm{ref}} - \tau_{\mathrm{s}}(k))^2,$$

which is zero when the desired sawtooth period is equal to the actual sawtooth period, and greater than zero otherwise. In practice, it is often encountered that the cost function is shaped to get a desired convergence rate throughout the working range. However, this implicitly includes knowledge on the DC-gain of the process in the cost function. Variations on the plasma parameters can significantly change the I/O map as shown in Fig. 2. Therefore, to guarantee high robustness, the cost function is not shaped.

B. Gradient estimator design

The task of the gradient estimator is to determine the derivatives of the cost function with respect to the nominal operating point $\hat{\vartheta}$. For the controller proposed in this paper, only the first-order gradient is needed. The working principle of the gradient estimator relies on forcing a perturbation on the input of the process. Consider the combination of the sawtooth process and the cost function in a static algebraic function $f(\vartheta(k))$. The sawtooth period can be approximated as a static function of the launcher angle if the controller operates sufficiently slow. A schematic representation of a minimal gradient estimator as proposed in [27] is shown in Fig. 4. For its output $\xi(k)$ it follows that $\xi(k) = d(k)f(\hat{\vartheta} + d(k))$. Here the perturbation is chosen to be sinusoidal. This is a common choice in ESC, although other perturbations could in principle also be used [28]. Let ω be the frequency and α the amplitude of the perturbation, so $d(k) = \alpha\sin(\omega k)$, we then obtain

$$\xi(k) = \alpha\sin(\omega k) f(\hat{\vartheta} + \alpha\sin(\omega k)). \tag{1}$$

This equation can be approximated by using a first-order Taylor expansion of $f(\vartheta(k))$ around the nominal input $\hat{\vartheta}$,

$$f(\hat{\vartheta}+\alpha\sin(\omega k)) \approx f(\hat{\vartheta}) + \alpha\sin(\omega k)\left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\vartheta=\hat{\vartheta}} + \mathscr{O}(\alpha^2). \tag{2}$$

Since α is typically chosen small, the higher-order terms in α are neglected. Substitution of this result in (1) yields

$$\xi(k) \approx \alpha\sin(\omega k) f(\hat{\vartheta}) + \alpha^2\sin^2(\omega k)\left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\vartheta=\hat{\vartheta}} = \alpha\sin(\omega k) f(\hat{\vartheta}) + \frac{\alpha^2}{2}(1-\cos(2\omega k))\left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\vartheta=\hat{\vartheta}}. \tag{3}$$

This shows that $\xi(k)$ consists of a static gradient-dependent component with additional oscillations due to the perturbation. The optimizer, which will use this $\xi(k)$, operates on a longer time-scale than the estimator; the oscillations will effectively average out over such a long time-frame, since

$$\lim_{K\to\infty}\frac{1}{K}\sum_{k=0}^{K}(\xi(k)) \approx \frac{\alpha^2}{2}\left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\vartheta=\hat{\vartheta}}. \tag{4}$$

Hence, under the assumption of time-scale separation, the output $\xi(k)$ of Fig. 4 indeed provides an estimation of the gradient of the cost function (in an averaged sense) with a scaling of $2\alpha^{-2}$.

The accuracy of the gradient estimator is further improved by the inclusion of additional filters. In the complete control system, the nominal operating point $\hat{\vartheta}$ is time-varying. Consequently, the term $f(\hat{\vartheta})$ in (2) will not average out completely. Since $\hat{\vartheta}$ is typically varying slowly, a high-pass filter can be applied right after the cost function to attenuate $f(\hat{\vartheta})$ without removing the response to the perturbation, i.e., the term $\alpha\sin(\omega k)\left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\vartheta=\hat{\vartheta}}$ in (2). The filtered signal $y'(k)$ must have an average close to zero, while keeping as much of the original frequency content of $y(k)$. The inclusion of this filter improves the gradient estimation, but ξ is still an approximation of the DC-gain which is only valid in an averaged sense.

Additionally, a low-pass filter could be applied on $\xi(k)$ [17], [18], [21]. If this filter is tuned to have a slow response, the oscillating terms in (3) are attenuated at the cost of delay. A more elegant approach is the application of a moving average filter. Note from (3) that $\xi(k)$ consists of sums of periodic signals of frequency $n\omega$ with $n = 1,2,3,\ldots$. A moving average filter, with the time window equal to the period time of the perturbation, suppresses the exact same frequencies, thereby removing all oscillating terms in (3) completely. Let the perturbation frequency be $\omega = a\pi$ where $0 < a < 1$ and $2/a$ must be a natural number. The perturbation then becomes

$$d(k) = \alpha\sin(a\pi k), \tag{5}$$

and a full period of the perturbation then involves $n = 2/a$ sawtooth crashes. Therefore, we propose the following moving average filter: $\xi(k) = \frac{1}{n}\sum_{j=1+k-n}^{k}\xi'(j)$, with $\xi'(k)$ the unfiltered gradient estimate. The output $\xi(k)$ is the DC-gain estimate, again accurate up to scaling with $2\alpha^{-2}$, see (4). The final design of the gradient estimator is shown in Fig. 5 and indicated with the dashed gray box.

C. Optimizer design

The task of the optimizer is to change the launcher angle such that the cost function f is minimized. Hereto, gradient information of the cost function is required, which is provided by the gradient estimator.

To show that the gradient is necessary for the optimization, consider a commonly used optimizer [17], [18], [27]:

$$\vartheta(k+1) = \vartheta(k) - \gamma\frac{\mathrm{d}f}{\mathrm{d}\vartheta}(k). \tag{6}$$

The optimization rate is proportional to the gradient of the cost function, and tuned with $\gamma > 0$. This optimizer is known as first-order gradient descent. If the gradient is positive, the estimate of the minimizer ϑ is updated in the negative direction, and vice versa, until the minimum is reached.

The gain γ scales the convergence rate of the optimizer, which depends on the estimated gradient. This implies that for processes with large DC-gain, γ has to be tuned small to ensure that the optimization is sufficiently slow. Unfortunately, for the sawtooth control problem, the system undergoes very large changes in DC-gain as shown in Fig. 2. If γ is tuned such that acceptable performance is achieved at the regions where the system has large gain, the convergence speed is extremely slow in the regions with small DC-gain. If the optimizer tuned such that it is performing well in the region with small DC-gain it is actually unstable in the region where the DC-gain is large, since the timescale separation is not guaranteed. This issue was encountered in [16] and ameliorated by scheduling the optimizer gain which compromises robustness since this requires information on the process.

An alternative approach towards tackling this problem is to make the convergence rate completely independent of the DC-gain by means of a so-called sliding mode optimizer [29]. The idea is to only use the sign of the first-order gradient estimate, and steer the launcher angle with constant velocity (in discrete-time domain) to the minimizer. The final design of the optimizer is shown in Fig. 5

For this scheme, the estimate of the minimizer $\hat{\vartheta}(k)$ is updated with a constant rate, determined by γ, in °/crash towards the minimum. There are three main advantages of this type of optimizer:

- The DC-gain does not affect the convergence rate, hence acceptable performance can be achieved throughout the entire feasible operating range for ϑ;
- Only the sign of the gradient is needed for optimization, hence the ESC is more robust;
- The tuning of the optimizer gain is simplified considerably.

There are also disadvantages:

- Chatter of ϑ around the optimum, a phenomenon explained in further detail below;
- Constant convergence rate; hence the rate is suboptimal for some regions of the feasible operating range.

Chatter is referred to as the unwanted bouncing or switching of variables, in this case the launcher angle ϑ. The proposed optimizer is always optimizing with a fixed convergence rate determined by γ. As a result, when close to an optimum, $\hat{\vartheta}(k)$ will 'bounce' around the optimum.

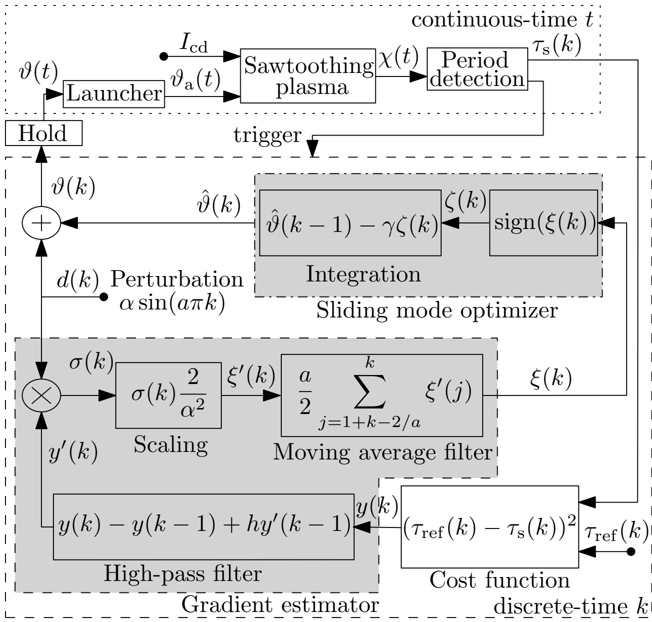

Fig. 5. Closed-loop system. The process runs in continuous time and is indicated with the dotted box. The controller runs in discrete time and is indicated with the dashed box, inside the controller the cost function is indicated with the white box, the gradient estimator is indicated with the dashed gray box, and the optimizer is indicated with the dash-dotted gray box.

IV. Tuning of the controller parameters

The perturbation frequency a should be selected such that it lies in the pass-band of the process, i.e. the resulting oscillation on the output of the process should appear with as little phase lag as possible, or with as little variation in phase lag as possible (e.g. induced by changes in operating conditions), since constant phase lag can be compensated for [19]. The frequency is selected $a = 0.2$ crash^{-1}. The perturbation amplitude is typically selected small, since large amplitudes introduce a large estimation error. The minimal amplitude is limited by the response of the mechanical launcher [30], and is therefore selected as $\alpha = 0.3°$. The cut-off frequency of the high-pass filter should be chosen such that the low frequency attenuation is as large as possible, but the perturbation frequency should lie in the pass-band of the filter; $h = 0.9$ is selected. The optimizer gain is more difficult to select; it should be sufficiently slow such that the gradient estimate can converge close enough to the actual gradient of the cost function. By iteratively increasing the optimizer gain a value of $\gamma = 0.02°$/crash is conceived.

V. Simulation results

In the simulation results, the response of the launcher to a desired angle is assumed to be ideal, i.e. $\vartheta_{\mathrm{a}} = \vartheta$. The controller's performance is tested for a variation in plasma parameters and for added disturbances on the current drive I_{CD}, actual launcher angle ϑ_{a} and sawtooth period τ_{s}.

A. Robustness against varying plasma parameters

The robustness of the controller is tested for a time-varying reference profile τ_{ref} and a change in plasma parameters, Fig. 6 depicts the results. Fig. 6c shows that the steady-state I/O map has an entirely different shape for the two different plasma

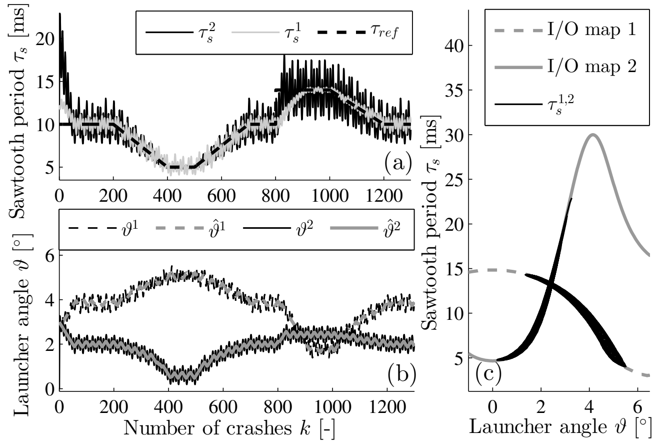

Fig. 6. Simulation results with different plasma parameters, the original simulation is indicated with (^{1}) and (^{2}) is with the changed plasma parameters. The sawtooth period τ_{s} is shown in (a), the launcher angle ϑ and estimated minimizer $\hat{\vartheta}$ in (b) and the sawtooth period trajectory in (c).

parameter settings. Nevertheless, the reference is successfully tracked for both cases, see Fig. 6a. The oscillation on the sawtooth period for simulation 2 at $\tau_s \approx 15$ ms is larger since the DC-gain at that operating point is larger. At $\tau_s \approx 5$ ms the oscillation is smaller, since that specific operating point has a DC-gain close to zero. The settling time after the step at 800 crashes is much faster for simulation 2, which is also caused by the larger DC-gain at $\tau_s \approx 15$.

A linear controller as used in [12], [13], [15] is not able to cope with disturbances or process variations which change the sign of the DC-gain. In the case presented in Fig. 6, it has in fact changed over a large portion of the operating range. Note that the change of plasma parameters is made off-line. An on-line change is also possible; this is similar to changing the reference sawtooth period. Since the controller is able to handle stepwise changes in the reference, it is expected that it can handle stepwise changes of the plasma parameters as well; both introduce a sudden change in the cost function.

B. Robustness against noise and sawtooth detection delay

The robustness of the proposed control strategy against sawtooth crash detection delay and noise on the in- and outputs is also tested in simulation. A physical example of noise on the launcher angle would be the interaction of the ECCD beam with the boundary of the plasma. Turbulent effects in the edge of the plasma, which occur on a very short timescale, can deflect the ECCD beam. The current deposition is thus 'smeared' over a small region instead of a fixed location. Noise in the current drive could be caused by vibrations in the gyrotron due to cooling systems and the output noise on the sawtooth period can be caused by small errors in the period detection. Furthermore, a sawtooth crash detection delay of 5 ms is added. A crash detection delay increases the phase shift between the perturbation and the resulting oscillation on the measured sawtooth period; hence the gradient estimate is influenced. The simulation results are depicted in Fig. 7. It shows that the controller is successfully tracking the reference and is thus robust for the added disturbances. With a detection delay of 5 ms, the introduced phase shift in the oscillation on the output is at

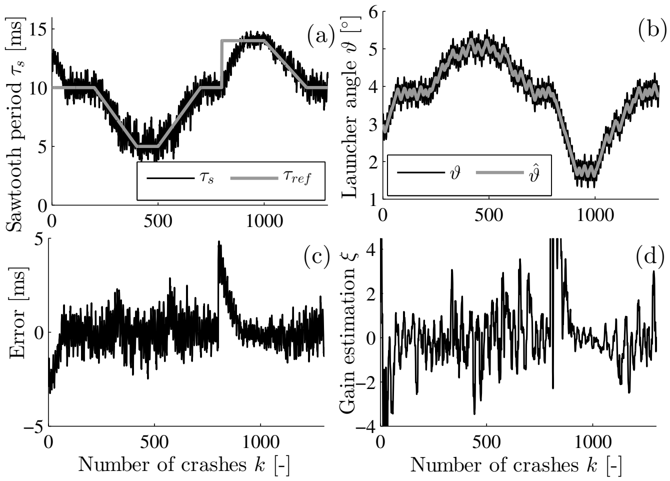

Fig. 7. Simulation results with disturbances present: (a) the sawtooth period τ_s, (b) the launcher angle ϑ and estimated minimizer $\hat{\vartheta}$, (c) the tracking error $\tau_s - \tau_{ref}$ and (d) the DC-gain estimate ξ.

most 1 crash. Since the chosen perturbation period is 10 crashes this is only a minor influence. If the perturbation is slow enough, the moving average filter attenuates the high-frequent noise from the gradient estimate.

VI. Conclusions and Future Work

This research provides a structured design of a robust sawtooth period controller. A sawtooth model was used to benchmark the performance of the controller, although the controller is applicable to more comprehensive models and experiments as well.

The working principle of the proposed Extremum Seeking Controller (ESC) relies on on-line identification of the gradient of a cost function by a gradient estimator. The ESC operates on a slower timescale that the sawtooth period dynamics. The benefit of this control strategy is that very little information of the sawtooth instability is required. The topology of the proposed controller is designed such that high robustness with acceptable performance is achieved. The controller does not rely on any model describing the behaviour of the sawtooth period and is therefore very robust.

Simulation results showed the tracking of sawtooth period references under different conditions. The steady-state tracking error is determined by the perturbation employed in the ESC scheme and the chatter introduced by the sliding mode optimizer used. The controller is able to handle step-wise changes in the sawtooth period reference, changes in plasma parameters, crash detection delay and noise on the in- and outputs of the sawtooth process.

Both the high robustness and the small amount of process knowledge required make ESC a very interesting candidate to apply in practice. The possible applications of the proposed ESC are not limited to the sawtooth control problem only, the controller could be an interesting candidate strategy for a wide range of fusion-related control problems with high model uncertainty.

References

[1] K. Schultz, "Why fusion? a discussion of energy alternatives," *IEEE Control Systems Magazine*, vol. 26, pp. 32–34, 2006.

[2] J. A. Wesson and D. J. Campbell, *Tokamaks*. Oxford university press, 2004.

[3] R. J. Hastie, "Sawtooth instability in tokamak plasmas," *Astrophysics and Space Science*, vol. 256, pp. 177–204, 1998.

[4] S. von Goeler, W. Stodiek, *et al.*, "Studies of internal disruptions and m=1 oscillations in tokamak discharges with soft x-ray techniques," *Phys. Rev. Letters*, vol. 33, pp. 1201–1203, 1974.

[5] M. F. F. Nave, J. Rapp, *et al.*, "Role of sawtooth in avoiding impurity discharges," *Nuclear Fusion*, vol. 43, pp. 1204–1213, 2003.

[6] O. Sauter, E. Westerhof, *et al.*, "Control of neoclassical tearing modes by sawtooth control," *Phys. Rev. Letters*, vol. 88, 2002.

[7] A. Gude, S. Günther, *et al.*, "Temporal evolution of neoclassical tearing modes and its effect on confinement reduction in ASDEX Upgrade," *Nuclear Fusion*, vol. 42, pp. 833–840, 2002.

[8] A. Mück, T. P. Goodman, *et al.*, "Sawtooth control experiments on asdex upgrade," *Plasma Physics and Controlled Fusion*, vol. 45, pp. 1633–1655, 2005.

[9] O. Merkulov, *Control of plasma profiles and stability through localised Electron Cyclotron Current Drive*, 2006.

[10] C. Angioni, T. P. Goodman, *et al.*, "Effects of localized electron heating and current drive on the sawtooth period," *Nuclear Fusion*, vol. 43, pp. 455–465, 2003.

[11] M. van Berkel, G. Witvoet, *et al.*, "Real-time wavelet detection of crashes in limit cycles of non-stationary fusion plasmas," *Fusion Engineering and Design*, vol. 86, pp. 2908–2919, 2011, in press.

[12] J. I. Paley, F. Felici, *et al.*, "Real time control of the sawtooth period using EC launchers," *Plasma Physics and Controlled Fusion*, vol. 51, pp. 1–9, 2009.

[13] M. Lennholm, L.-G. Eriksson, *et al.*, "Closed loop sawtooth period control using variable ECCD injection angles on tore supra," *Fusion Science and Technology*, vol. 55, pp. 45–55, 2009.

[14] ——, "Demonstration of effective control of fast-ion-stabilized sawteeth by electron-cyclotron current drive," *Phys. Rev. Letters*, vol. 102, 2009.

[15] G. Witvoet, M. R. de Baar, *et al.*, "Systematic design of a sawtooth period feedback control using a Kadomtsev–Porcelli sawtooth model," *Nuclear Fusion*, vol. 51, no. 7, p. 073024.

[16] J. Paley, F. Felici, *et al.*, "From profile to sawtooth control: developing feedback control using ECRH/ECCD systems on the TCV tokamak," *Plasma Physics and Controlled Fusion*, vol. 51, pp. 1–11, 2009.

[17] Y. Tan, W. H. Moase, *et al.*, "Extremum seeking from 1922 to 2010," *Proc. of the 29th Chinese Control Conf., Beijing, China*, pp. 14–26, 2010.

[18] M. Krstić and H.-H. Wang, "Stability of extremum seeking feedback for general nonlinear dynamic systems," *Automatica*, vol. 36, pp. 595–601, 2000.

[19] M. Krstić, "Performance improvement and limitations in extremum seeking control," *Systems & Control Letters*, vol. 39, pp. 313–326, 2000.

[20] D. Nešić, "Extremum seeking control: convergence analysis," *European Journal of Control*, vol. 15, no. 3-4, pp. 331–347, 2009.

[21] D. Nešić, Y. Tan, *et al.*, "A unifying approach to extremum seeking: adaptive schemes based on estimation of derivatives," *49th IEEE Conf. on Decision and Control, Atlanta, Georgia, USA*, pp. 4625–4630, 2010.

[22] G. Witvoet, M. Lauret, *et al.*, "Numerical demonstration of injection locking of the sawtooth period by means of modulated EC current drive," *Nuclear Fusion*, vol. 51, no. 10, p. 103043, 2011.

[23] G. Witvoet, M. Steinbuch, *et al.*, "Sawtooth period control strategies and designs for improved performance," *Nuclear Fusion, in press*, 2012.

[24] M. Lauret, F. Felici, *et al.*, "Demonstration of sawtooth period locking in TCV plasmas," *Accepted for publication in Nuclear Fusion*, 2012.

[25] J.-Y. Choi, M. Krstić, *et al.*, "Extremum seeking for discrete-time systems," *IEEE Transactions on automatic control*, vol. 47, no. 2, pp. 318–323, 2002.

[26] J. Bolder, G.Witvoet, *et al.*, "Robust sawtooth period control based on adaptive online optimization," *Nuclear Fusion, in press*, 2012.

[27] Y. Tan, D. Nešić, *et al.*, "On non-local stability properties of extremum seeking control," *Automatica*, vol. 42, pp. 889–903, 2006.

[28] ——, "On the choice of dither in extremum seeking systems: a case study," *Automatica*, vol. 44, pp. 1446–1450, 2008.

[29] Y. Pan, U. Özgüner, *et al.*, "Stability and performance improvement of extremum seeking control with sliding mode," *International Journal of Control*, vol. 76, pp. 968–985, 2003.

[30] B. A. Hennen, E. Westerhof, *et al.*, "A closed-loop control system for stabilization of MHD events on TEXTOR," *Fusion Engineering and Design*, vol. 84, pp. 928–934, 2009.