Robust sawtooth period control based on adaptive online optimization

2012 Nucl. Fusion 52 074006

(http://iopscience.iop.org/0029-5515/52/7/074006)

Nucl. Fusion 52 (2012) 074006 (16pp) doi:10.1088/0029-5515/52/7/074006

Robust sawtooth period control based on adaptive online optimization

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]

[1] Eindhoven University of Technology, Department of Mechanical Engineering, Control Systems Technology Group, PO Box 513, 5600 MB Eindhoven, The Netherlands
[2] Eindhoven University of Technology, Department of Mechanical Engineering, Dynamics and Control Group, PO Box 513, 5600 MB Eindhoven, The Netherlands
[3] FOM-Institute for Plasma Physics Rijnhuizen, Association EURATOM-FOM, Trilateral Euregio Cluster, PO Box 1207, 3430 BE Nieuwegein, The Netherlands
[4] TNO Technical Sciences, Opto-mechatronics, PO Box 155, 2600 AD Delft, The Netherlands

E-mail: gert.witvoet@tno.nl

Received 8 September 2011, accepted for publication 17 February 2012
Published 5 July 2012
Online at stacks.iop.org/NF/52/074006

Abstract
The systematic design of a robust adaptive control strategy for the sawtooth period using electron cyclotron current drive (ECCD) is presented. Recent developments in extremum seeking control (ESC) are employed to derive an optimized controller structure and offer practical tuning guidelines for its parameters. In this technique a cost function in terms of the desired sawtooth period is optimized online by changing the ECCD deposition location based on online estimations of the gradient of the cost function. The controller design does not require a detailed model of the sawtooth instability. Therefore, the proposed ESC is widely applicable to any sawtoothing plasma or plasma simulation and is inherently robust against uncertainties or plasma variations. Moreover, it can handle a broad class of disturbances. This is demonstrated by time-domain simulations, which show successful tracking of time-varying sawtooth period references throughout the whole operating space, even in the presence of variations in plasma parameters, disturbances and slow launcher mirror dynamics. Due to its simplicity and robustness the proposed ESC is a valuable sawtooth control candidate for any experimental tokamak plasma, and may even be applicable to other fusion-related control problems.

1. Introduction

The sawtooth instability is a periodic redistribution of the plasma core particles and energy [1–3]. On the one hand, its mixing effect provides a mechanism to regulate the exhaust of helium ash and α-particles [4], and the influx of deuterium and tritium in a fusion reactor. On the other hand, the sawtooth instability can trigger neo-classical tearing modes (NTMs), which in turn reduce the operational performance and could lead to disruptions [5, 6]. Sawtooth control, in particular control of the sawtooth period [7, 8], is necessary to avoid NTM triggering while concurrently refreshing the plasma core.

The onset of a sawtooth crash is often associated with the magnetic shear at the $q = 1$ surface [9]. The sawtooth period can therefore be affected by changing the shear around $q = 1$ through the injection of electron cyclotron (EC) waves, see [10–13] and references therein. By changing the deposition location of the resulting EC current drive (ECCD) relative to the $q = 1$ surface the growth rate of the shear can either be increased or decreased, leading to shorter or longer sawtooth periods, respectively. The ECCD deposition location is typically determined by the angle of an EC mirror. Sawteeth are usually observed using soft x-ray or electron cyclotron emission (ECE) measurements. The sawtooth period can be extracted from these measurements in several ways, e.g. using multiresolution wavelet analysis as discussed in [14].

The ECCD actuator has successfully been employed in a closed loop to control the sawtooth period on both TCV [15] and Tore Supra [16, 17] using classical linear controllers. A systematic design for such controllers has been presented in [18] based on structured analysis of the dynamics of the sawtooth period. This approach enabled *a priori* assessment of closed-loop stability, performance and robustness, which has been verified by simulations. Recently a similar methodology has been used to suggest performance improvement strategies [19], yielding faster convergence of the sawtooth period with high accuracy.

The above controller design strategies assume that the dynamics of the sawtooth period, i.e. its frequency response functions [18], and its statics, i.e. its steady-state input–output map, can be measured and remain approximately constant over time. This asks for reproducible and predictable plasmas. Under reactor-like conditions these requirements can easily be met, but under experimental circumstances plasma parameters are often uncertain or unknown. Variations in the profiles of, e.g., density, electron temperature, conductivity and impurity concentration can cause the $q = 1$ surface to shift, and effects like scattering and EC beam deflections can alter the ECCD deposition location and width. Consequently, the input–output map and the underlying dynamics can change significantly, due to which the above mentioned controllers can become unstable. Hence, experimental devices require a sawtooth period controller which is robust against large variations, uncertainties and disturbances.

A first example of such a robust alternative has been presented in [20], showing experimental results from TCV using extremum seeking control [21]. The work in [20] is based on the specific controller structure from [22, 23], and maximizes the sawtooth period without *a priori* knowledge of the ECCD deposition location which yields this maximum period. However, the controller parameters were reported to be difficult to select, and stabilization on smaller-than-maximum periods was not addressed.

Recent developments in control engineering [21, 24, 25] have shed new light on extremum seeking control (ESC), and have generalized and extended the work in [22, 23]. ESC is essentially an adaptive control strategy, which optimizes a certain cost function using an online estimation of its gradient. The results in [21, 24, 25] have shown that the ESC structure can be decomposed into separate subsystems, i.e. a cost function, a gradient estimator and an optimization routine. Each subsystem can be chosen and designed separately, which offers flexibility and allows for a systematic design and tuning of the extremum seeking controller. These insights thus offer solutions for the practical issues raised in [20], and allow for a dedicated optimization of ESC for the sawtooth period control problem.

In this paper we present such an extremum seeker design, tailored to the sawtooth period control problem based on these new developments [21, 25]. The controller is defined in the crash-driven discrete-time framework presented in [18] and uses a cost function which allows tracking of any desired sawtooth period. The gradient of this function is estimated online by means of an external perturbation on the EC mirror angle and subsequent dedicated filtering. A so-called sliding mode optimizer [26] uses this gradient to adjust the mirror angle with a constant rate towards the optimum, which coincides with the desired sawtooth period. Practical guidelines will be provided to tune the controller parameters, based on the required separation of time-scales between these subsystems of the extremum seeker [21].

This proposed adaptive control strategy is model-free, i.e. it does not rely on any mathematical or physical model of the sawtooth instability. Therefore, the controller is inherently robust, as it can be applied to any sawtooth simulation or sawtoothing plasma. A Kadomtsev–Porcelli sawtooth model [18] is used as a case study to benchmark and test the controller. Simulation results affirm that the controller can track any time-varying sawtooth period reference, even when the plasma parameters change significantly, various types of disturbances and a detection delay are added, or slow EC launcher dynamics is incorporated in the control loop. This demonstrates that ESC is indeed highly robust against uncertainties and disturbances, and is therefore readily applicable on experimental devices to control the sawtooth period. The high robustness does come at the expense of degraded closed-loop performance compared with the techniques in [18, 19]; the separation of time-scales prescribes low convergence speed and the external perturbation induces ongoing oscillations on the sawtooth period.

This paper is organized as follows, section 2 briefly discusses the considered sawtooth model and the sawtooth period control problem. In section 3, the basic principles of extremum seeking for sawtooth control are discussed. Controller structure design and tuning guidelines are elaborated in section 4. Tests of the controller on the sawtooth model are presented in section 5 and possible performance improvements are suggested in section 6. Conclusions and discussion are addressed in section 7.

2. Control problem formulation

Since the highly robust control strategy for the sawtooth period presented in this paper is model-free, it is applicable to any sawtooth model or experimental sawtoothing plasma. However, to illustrate the tuning of its controller parameters and to assess the resulting closed-loop behaviour, ESC will be applied to the specific control-oriented sawtooth model proposed in [18]. In this section this model is briefly recapitulated, and used to introduce the sawtooth period control problem.

2.1. Kadomtsev–Porcelli sawtooth model

The considered sawtooth model consists of three main elements: the diffusion equation of the poloidal magnetic field, Porcelli's criterion for triggering of the sawtooth crash [9] and the Kadomtsev reconnection model [27]. The model consists of a set of equations to describe the evolution of the poloidal magnetic field $B_\theta(r, t)$ as a function of time t and radius r

$$\frac{\partial}{\partial t}B_\theta = \frac{\partial}{\partial r}\left(\frac{\eta}{\mu_0 r}\left(B_\theta + r\frac{\partial}{\partial r}B_\theta\right) - \eta J_{\mathrm{CD}}\right) \quad \text{if } s_1 < s_{\mathrm{crit}}, \tag{1a}$$

$$B_\theta(r, t^+) = \begin{cases} B_\theta(r, t^-) & \text{for } r \geqslant r_{\mathrm{mix}} \\ \frac{\mathrm{d}}{\mathrm{d}r}\Psi_*^c(r) + \frac{1}{R_0} r B_\phi & \text{for } r < r_{\mathrm{mix}} \end{cases} \quad \text{if } s_1 \geqslant s_{\mathrm{crit}}, \tag{1b}$$

with boundary conditions $B_\theta(0, t) = 0$ and $B_\theta(a, t) = \mu_0 I_{\mathrm{p}}/2\pi a$. Here η is the plasma resistivity, μ_0 the magnetic permeability, R_0 the tokamak major radius, a its minor radius, B_ϕ the toroidal magnetic field and I_{p} the plasma current. Moreover, $s_1 = s(r_{q=1})$ is the magnetic shear at the surface where the safety factor q equals unity, where for large aspect ratio tokamaks $q(r, t) = rB_\phi/R_0 B_\theta$ and $s(r, t) = \frac{r}{q}\frac{\mathrm{d}q}{\mathrm{d}r}$.

According to Porcelli's model [9] a sawtooth crash is triggered if this s_1 exceeds a critical value $s_{\rm crit}$, which is assumed to be constant here. At a sawtooth crash, the flux surfaces up to the mixing radius $r_{\rm mix}$ reconnect on a very short time-scale according to the model of Kadomtsev [27], represented by equation (1*b*). The magnetic field B_θ after a crash (at time t^+) follows from the post-crash helical flux function $\Psi_*^c(r)$, which depends on the pre-crash helical flux function $\Psi_*(r)$ at time t^-. This $\Psi_*^c(r)$ is calculated using the Archimedes–Kadomtsev approach proposed in [18], i.e.

$$\Psi_*^c(r_c) = \Psi_*(r_{i-}) = \Psi_*(r_{j+}), \quad \text{where} \quad r_c^2 = \sum_i r_{i-}^2 - \sum_j r_{j+}^2, \tag{2}$$

where r_{i-} denotes the pre-crash surfaces with $\mathrm{d}\Psi_*/\mathrm{d}r < 0$, and r_{j+} the pre-crash surfaces with $\mathrm{d}\Psi_*/\mathrm{d}r > 0$. The output of the model is the time between two subsequent crashes equation (1*b*), which defines the sawtooth period $\tau_{\rm s}$.

The magnetic field B_θ, and thereby the period $\tau_{\rm s}$, can be influenced by the EC current drive profile $J_{\rm CD}$ in equation (1*a*). This profile is determined by the total driven current $I_{\rm CD}$ and the deposition location of the EC beam (assuming a constant deposition width). The latter is directly influenced by the EC mirror angle $\vartheta_{\rm a}$, which is considered the input of the sawtooth model. The expression to determine $J_{\rm CD}$ from the current drive $I_{\rm CD}$ and the EC mirror angle $\vartheta_{\rm a}$ is given in [18].

This relatively simple model is control-oriented in the sense that it captures the qualitative input–output behaviour of the sawtooth and its dominant (magnetic) time-scales [18]. Note that the model is only valid for circular tokamaks, does not incorporate heating effects of the EC beam and therefore does not encompass the hysteresis discussed in [15], and note that in reality $s_{\rm crit}$ is not constant. However, these issues merely affect the quantitative details of the input–output behaviour and are therefore irrelevant for the controller structure. In particular, the adaptive controller discussed in this paper is model-free and is thus explicitly robust against such modelling errors.

The model is implemented in a Matlab® Simulink® environment [18]. This gives great flexibility in the design and interconnection of systems and signals. The model parameters have been chosen according to the specifications of the TEXTOR tokamak [28], and tuned to yield a realistic ohmic sawtooth period of about 15 ms.

2.2. *The sawtooth period control problem*

A schematic representation of the sawtooth period control problem is shown in figure 1. The first input to the sawtoothing plasma, or in our case the sawtooth model, is the EC driven current $I_{\rm CD}$, which is in the same direction as the plasma current and is kept constant at $I_{\rm CD} = 2$ kA. The second input is the mirror angle $\vartheta_{\rm a}$; the requested mirror angle is denoted by ϑ, which is the control variable in this paper. The output of the sawtooth system is a set of measurements χ, e.g. soft x-ray or ECE, from which a sawtooth period $\tau_{\rm s}$ is determined by a period detection algorithm [14]. The controller has two inputs: $\tau_{\rm s}$ and a reference sawtooth period $\tau_{\rm s,ref}$, which may be time varying. The task of the controller is to steer the mirror such that $\tau_{\rm s}$ converges to $\tau_{\rm s,ref}$.

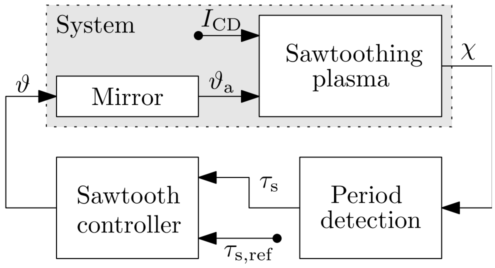

Figure 1. Sawtooth period controller topology, where χ represents a sawtoothing plasma diagnostic (such as ECE or soft x-ray).

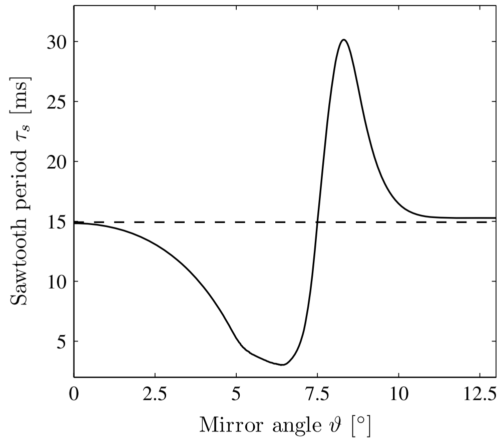

Figure 2. Steady-state input–output map of the model presented in section 2.1 (solid) and the ohmic sawtooth period τ_Ω (dashed), for $B_\phi = 2.45$ T and $I_{\rm p} = 400$ kA.

As we will later argue, our proposed control strategy operates on a time-scale that is slower than the time-scale of the sawtooth dynamics. This implies that the mirror adjustments are so slow that the sawtooth period is essentially always close to its steady-state value. For this reason the steady-state input–output behaviour of the sawtooth instability is of importance, which describes the relation between the mirror angle and the sawtooth period in steady-state, i.e. it depicts $\tau_{\rm s}$ for a certain ϑ as $t \to \infty$. Figure 2 shows this input–output map for the sawtooth model described in section 2.1, which we will use as a case study throughout this paper. The ohmic period τ_Ω is indicated by the dashed line. Figure 2 thus shows that for mirror angles below 7.5° the sawtooth period shortens, and for larger angles the period is lengthened. This corresponds to injecting current either inside or outside the $q = 1$ surface. Such behaviour is in agreement with previous observations [12, 13] and can be expected on ITER also [10]. Note that the slope or gradient $\mathrm{d}\tau_{\rm s}/\mathrm{d}\vartheta$ of the input–output map is not constant and even changes sign; for some values of ϑ an increase in mirror angle yields an increase of the period, whereas for other values of ϑ the period decreases. This shows the importance of the gradient for a controller design; in order to steer the mirror in the right direction, the sign of $\mathrm{d}\tau_{\rm s}/\mathrm{d}\vartheta$ has to be known.

Standard feedback control strategies as used in [15–18] or high-performance controllers as in [19] require the sign of the gradient to be constant, which limits their operating space, e.g. to $0° \leqslant \vartheta \leqslant 6°$. Moreover, they rely on knowledge of the gradient (also known as the DC-gain) and the underlying

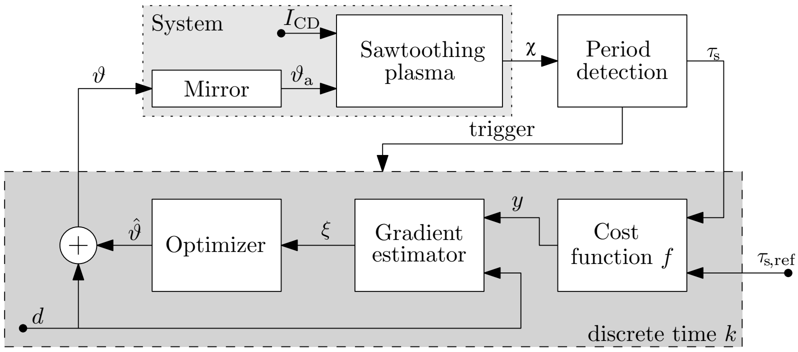

Figure 3. ESC topology for the sawtooth period, depicting the extremum seeking controller in the bottom grey box.

dynamics to ensure stability of the control loop. As mentioned in the introduction, this is a viable assumption under reactor-like conditions. However, on experimental tokamaks plasma parameters are often uncertain or time varying, leading to large variations in the sawtooth dynamics, while the controller operating space generally needs to be large. Hence, such devices require more robust control approaches.

A good candidate is extremum seeking control, which does not rely on any *a priori* knowledge of the system. It only requires the existence of a stable steady-state input–output map, with a unique output τ_s for each fixed input ϑ [22, 29]; the curve in figure 2, of which each point is indeed stable [18], indicates that the considered sawtooth model meets this criterion. ESC identifies the gradient of the system online and is therefore able to cope with uncertainties of the input–output map. Consequently, it can be used throughout the entire operating space of the mirror. In this paper we will discuss the systematic design of an extremum seeking (ES) controller for the sawtooth period where the robustness against variations of the input–output map is specifically addressed.

3. Fundamentals of extremum seeking for sawtooth period control

ESC is an adaptive control strategy that uses online optimization techniques to slowly drive a process to a desired operating point which minimizes a cost function f. For sawtooth period control this implies that ESC finds the mirror angle ϑ that minimizes a certain function f such that the period τ_s matches a desired reference value.

The block scheme in figure 3 shows the ESC topology for the sawtooth period control problem. This closed loop operates in both continuous time (sawtooth process) and discrete time (controller). The sawtooth period, i.e. the variable to be controlled, only changes when a crash occurs. There are no measurements of the period between crashes, hence, controller updates of the angle ϑ are triggered by the period detection. Therefore, the controller operates in discrete time [30], with the interval between control updates being the most recent sawtooth period, which concurs with the approach taken in [18]. The variable k is the crash counter and is a measure of discrete time.

The ES controller consists of three subsystems: a cost function, a gradient estimator and an optimizer. The cost function should be such that its function value y is minimal if the measured sawtooth period τ_s is equal to the reference sawtooth period $\tau_{s,ref}$. The gradient estimator uses a perturbation signal d to estimate the gradients of the cost function with respect to the mirror angle ϑ. The optimizer uses this gradient estimate ξ to drive the estimate of the optimal mirror angle $\hat{\vartheta}$ to the minimizer ϑ^* of the cost function, which is typically unknown, under the assumption that the input–output map is not known exactly.

For this controller to work properly, it should be designed such that a separation of the time-scales that each subsystem operates in is obtained [25]. The cost function is static, and could thus be viewed as part of the sawtooth process. The estimation of the gradient can only be performed correctly if the dynamics of the sawtooth period has converged sufficiently close to the steady-state. The perturbation d used by the gradient estimator thus has to be slower than the slowest time-scale of the sawtooth dynamics. 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 longest time-scale. In summary, an ES control loop should display three important time-scales [21, 25]:

(1) the dynamics of the sawtooth period, including the cost function (fastest);
(3) the perturbation for the gradient estimation (intermediate);
(3) the optimization (slowest).

This time-scale separation can be visualized in figure 4, which is inspired by [21]. It shows an arbitrary steady-state map (combination of sawtooth system and cost function) in grey, with input ϑ and the value y of the cost function as output. The output quickly evolves from some initial condition to the steady-state map. The gradient estimation scans the cost function on an intermediate time-scale, and the optimization is performed on the slowest time-scale. The separation of time-scales allows us to design each subsystem separately, starting with the fastest system.

4. Extremum seeking controller design

This section discusses the working principles of the different ESC subsystems, as well as the accompanying design procedures. Moreover, implementation and tuning guidelines are proposed in section 4.4, and a supervisory control loop is discussed in section 4.5.

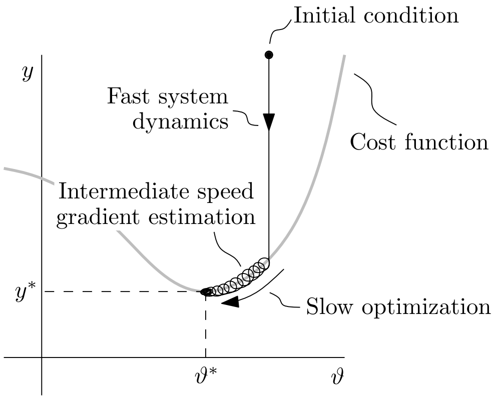

Figure 4. Output trajectory showing the time-scale separation.

4.1. The cost function

The task of the extremum seeker is to find a mirror angle ϑ that minimizes a static cost function f. This minimum of the cost function should correspond to a desired operating point, i.e. when the sawtooth period τ_s matches a desired period or reference $\tau_{s,ref}$. Therefore, we propose the cost function

$$f(\tau_s, \tau_{s,ref}) = (\tau_{s,ref} - \tau_s)^2, \tag{3}$$

which has a unique minimum at $\tau_s = \tau_{s,ref}$. Figure 5 depicts this cost function as a function of the input ϑ in steady-state for $\tau_{s,ref} = 10$ ms, obtained by substituting the input–output map of figure 2 into equation (3). It shows multiple optima that minimize the cost function, since there are two possible mirror angles that can yield a sawtooth period of 10 ms. It is unknown *a priori* to which minimum the extremum seeker will converge. In section 5.1 it will be shown that the dynamic behaviour of the closed-loop system is very different at the two optimal mirror angles. Moreover, note that in a specific case where $\tau_{s,ref} > \tau_\Omega$ there will also be a local non-optimal minimum of the cost function at $\vartheta = 0°$. In principle, ESC could converge to such a point. To avoid this behaviour a supervisor could be used, as discussed in section 4.5.

In practice the desired sawtooth period $\tau_{s,ref}$ may be time varying, due to which the cost function may change. However, as long as this reference is varied on a slower time-scale than the optimizer, the controller is able to track the moving optimum. In a previous application of ESC [20] the sawtooth period has been maximized. In our framework this can be achieved by choosing $\tau_{s,ref}$ in equation (3) greater than the maximum sawtooth period (i.e. >30 ms). Similarly, choosing $\tau_{s,ref} < 3$ ms will minimize the sawtooth period.

4.2. Gradient estimator design

Numerical optimization algorithms often make use of the gradient of a cost function to find its minimum. A commonly used optimizer [21, 22, 29] is the first-order gradient descent method

$$\hat{\vartheta}(k) = \hat{\vartheta}(k-1) - \gamma \cdot \left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\hat{\vartheta}(k-1)}, \tag{4}$$

where $\mathrm{d}f/\mathrm{d}\vartheta|_{\hat{\vartheta}(k)}$ is the gradient at crash k, i.e. the derivative of the cost function with respect to the control variable, in this case the mirror angle ϑ. The optimization rate is proportional

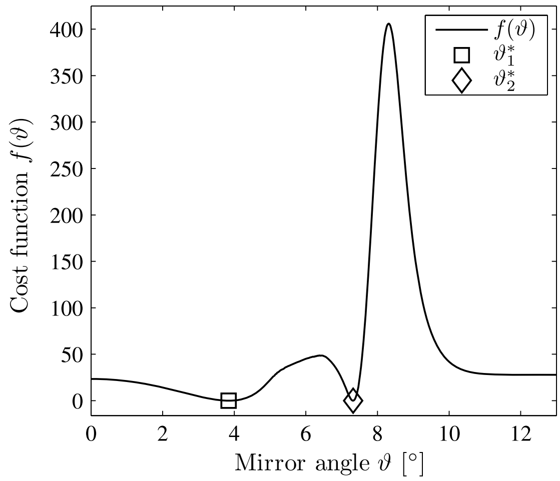

Figure 5. Steady-state cost function equation (3) evaluated for $\tau_{s,ref} = 10$ ms and the input–output map in figure 2. The optima ϑ_1^*, ϑ_2^* have been indicated with the markers.

to this gradient (scaled with a gain $\gamma > 0$); if the gradient is positive, the ϑ is decreased and vice versa, until the minimum is reached.

Since by assumption the input–output map is unknown, the gradient $\mathrm{d}f/\mathrm{d}\vartheta|_{\hat{\vartheta}(k)}$ has to be estimated online. To this end we first consider the minimal gradient estimator [29], depicted in grey in figure 6, whose working principle relies on using external perturbations on the input ϑ. In this scheme $y(k)$ is the value of the cost function, $\vartheta(k)$ its argument, $d(k)$ is an externally applied perturbation and $\xi(k)$ is the output variable. In accordance to the separation of time-scales we assume that the perturbation $d(k)$ is slower than the sawtooth dynamics, so that the combination of the sawtooth system and the cost function can be approximated by a static function $f(\vartheta(k))$, like the one shown in figure 5. For the output of the scheme it then follows that

$$\xi(k) = d(k) f(\hat{\vartheta} + d(k)). \tag{5}$$

The perturbation $d(k)$ is chosen to be sinusoidal. This is a common choice in ESC, although other perturbations are also possible [31]. Let ω be the frequency and α the amplitude, so that $d(k) = \alpha \sin(\omega k)$, then the output becomes

$$\xi(k) = \alpha \sin(\omega k) f(\hat{\vartheta} + \alpha \sin(\omega k)). \tag{6}$$

A first-order Taylor expansion of $f(\vartheta(k))$ around the nominal input $\hat{\vartheta}$ yields

$$f(\hat{\vartheta} + \alpha \sin(\omega k)) \approx f(\hat{\vartheta}) + \alpha \sin(\omega k) \left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\vartheta=\hat{\vartheta}} + \mathcal{O}(\alpha^2). \tag{7}$$

Assuming small α, and thereby neglecting the higher-order terms, substitution of this result in equation (6) then yields

$$\begin{aligned} \xi(k) &\approx \alpha \sin(\omega k) f(\hat{\vartheta}) + \alpha^2 \sin^2(\omega k) \left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\vartheta=\hat{\vartheta}} \\ &= \alpha \sin(\omega k) f(\hat{\vartheta}) + \frac{\alpha^2}{2}(1 - \cos(2\omega k)) \left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\vartheta=\hat{\vartheta}}. \end{aligned} \tag{8}$$

Hence, the instantaneous output $\xi(k)$ consists of a static component $\alpha^2/2$ times the gradient, plus additional oscillations

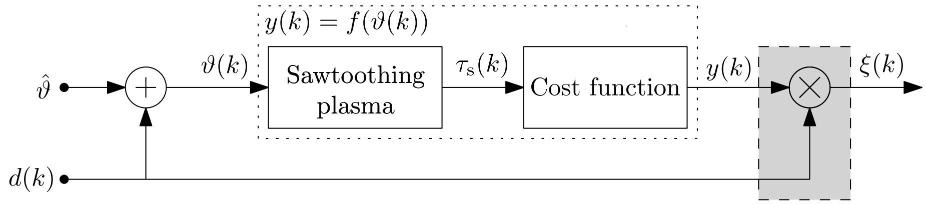

Figure 6. Topology of a minimal gradient estimator, indicated in grey.

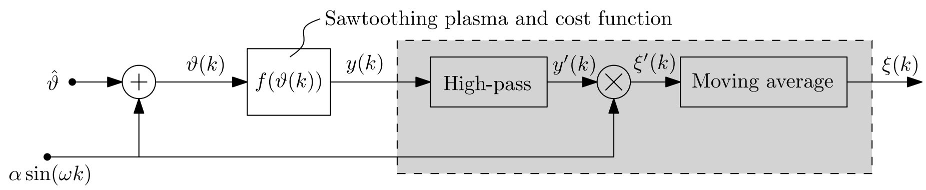

Figure 7. Block scheme of the proposed gradient estimator, indicated in grey.

due to the perturbation. However, note that 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

$$\lim_{K\to\infty} \frac{1}{K} \sum_{k=0}^{K} (\xi(k)) = \frac{\alpha^2}{2} \left.\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\right|_{\vartheta=\hat{\vartheta}}. \tag{9}$$

Hence, under the assumption of time-scale separation, the output $\xi(k)$ of figure 6 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 in equation (5) and figure 6 is further improved by the inclusion of additional filters. In the closed loop the nominal operating point $\hat{\vartheta}$ and thereby $f(\hat{\vartheta})$ are time varying. Consequently, the first term in equation (8) does not average out completely. Since $\hat{\vartheta}$ is varying slowly (with the speed of the optimizer), $f(\hat{\vartheta})$ can be attenuated by high-pass filtering the output of the cost function $y(k)$. The filtered signal $y'(k)$ then has an average close to zero, while the term $\alpha\sin(\omega k)\mathrm{d}f\mathrm{d}\vartheta|_{\vartheta=\hat{\vartheta}}$ in equation (7) is retained. The inclusion of this filter naturally improves the gradient estimation, although ξ is still an approximation which is only valid in an averaged sense.

Additionally, a low-pass filter could be applied to the gradient estimation [21, 22, 25]. If such a filter is tuned to have a slow response, the oscillating terms in equation (8) are attenuated at the cost of additional delay. A more elegant approach is the application of a moving average filter. Note that $\xi(k)$ consists of sums of periodic signals of frequency $n\omega$ with $n = 1, 2, 3, \ldots$. A moving average filter with a time window equal to the period time of the perturbation suppresses the exact same frequencies and thus removes all oscillating terms in equation (8). Let the perturbation frequency be $\omega = a\pi$ where $0 < a < 1$ and $2/a$ is a natural number. The perturbation then becomes

$$d(k) = \alpha\sin(a\pi k), \tag{10}$$

and a full period of the perturbation then involves $n = 2/a$ sawtooth crashes. Therefore, we propose the following moving average filter:

$$\xi(k) = \frac{1}{n} \sum_{j=1+k-n}^{k} \xi'(j), \tag{11}$$

with $\xi'(k)$ the unfiltered gradient estimate. A schematic representation of this extended gradient estimator, which is used in this paper, is shown in figure 7. The gradient estimate is formed by the output $\xi(k)$, which again has to be scaled by a factor $2\alpha^{-2}$.

The frequency of the perturbation separates the fast time-scale of the sawtooth dynamics from the slower time-scale of the gradient estimator. This frequency should be sufficiently low (belonging to the pass-band of the sawtooth dynamics [18]), such that the response of the sawtooth period to the perturbation is close to steady-state. On the other hand, a faster perturbation leads to faster convergence of the gradient estimator. The tuning of the perturbation frequency is discussed in section 4.4.

4.3. Optimizer design

The optimizer uses the estimation of the cost function gradient to change the mirror angle such that the cost function is minimized, such as the first-order gradient descent technique in equation (4). In this technique the gain γ scales the convergence rate of the optimizer, which depends on the estimated gradient. This implies that when the gradient is large, γ has to be small to ensure that the optimization is sufficiently slow. Unfortunately, in the sawtooth period control problem the gradient undergoes very large changes, as suggested by figures 2 and 5. If γ is tuned such that acceptable performance is achieved in the regions where the gradient is large, typically for mirror angles around 7.5°, the convergence speed is extremely slow in the regions with small gradients, typically for mirror angles $0° \leqslant \vartheta \leqslant 6°$. If the optimizer is tuned such that it is performing well in the region $0° \leqslant \vartheta \leqslant 6°$, it can actually be unstable in other regions, since the time-scale separation may not be guaranteed. This issue was encountered in [20] and ameliorated by scheduling the optimizer gain. However, this compromises robustness,

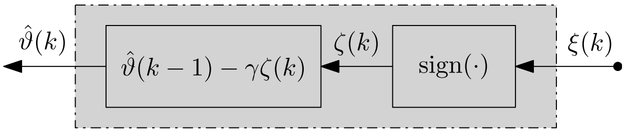

Figure 8. Sliding mode optimizer.

since it requires information on the input–output behaviour of the sawtooth.

An alternative approach is to make the convergence rate completely independent of the gradient by means of a so-called sliding mode optimizer [26]. This optimizer uses only the sign of the first-order gradient estimate, and steers the nominal mirror angle with constant velocity (in discrete time) towards the minimizer ϑ^*, i.e.

$$\hat{\vartheta}(k) = \hat{\vartheta}(k-1) - \gamma\,\text{sign}(\xi(k)). \tag{12}$$

The corresponding block scheme is shown in figure 8. There are three main advantages of this type of optimizer:

- the gradient does not affect the convergence rate, hence acceptable performance can be achieved throughout the entire operating range;
- it is more robust (with respect to convergence), since it requires only the sign of the gradient and thus allows for certain errors on the gradient estimation;
- the tuning of the optimizer gain is simplified considerably, as will be discussed in section 4.4.

A disadvantage of this approach is that chatter of ϑ around the optimum will arise. Chatter is referred to as the unwanted bouncing or switching of variables, in this case the mirror angle ϑ, as shown in figure 9. The proposed optimizer is always optimizing $\hat{\vartheta}(k)$ with a fixed convergence rate determined by γ. As a result, when close to an optimum, the mirror angle (and with that the sawtooth period) will display an additional oscillation around the optimum, with a lower frequency than the perturbation. This frequency is related to the settling time of the gradient estimate, and the slope of the trajectory of $\hat{\vartheta}(k)$ is equal to $\pm\gamma$. The chatter can be interpreted as a disturbance, but since its frequency is lower than the perturbation frequency, the correlation between perturbation and chatter is small. Hence, chatter does not influence the gradient estimate. The amplitude of the chatter is determined by the optimizer gain γ and the settling time of the optimizer, which is mainly determined by the perturbation frequency. Both a larger γ and a larger settling time of the optimizer yield a larger chatter amplitude.

4.4. Overall control system and tuning of the controller parameters

This section discusses the complete closed-loop control system shown in figure 10, and presents guidelines for the tuning of the controller parameters. Note again that the mirror and the plasma operate in continuous time, whereas the controller is defined in discrete time. The transition from continuous to discrete time is done by the sawtooth period detection. Conversely, a ‘hold’ operation keeps $\vartheta(k)$ constant in between crashes to form the continuous-time signal $\vartheta(t)$.

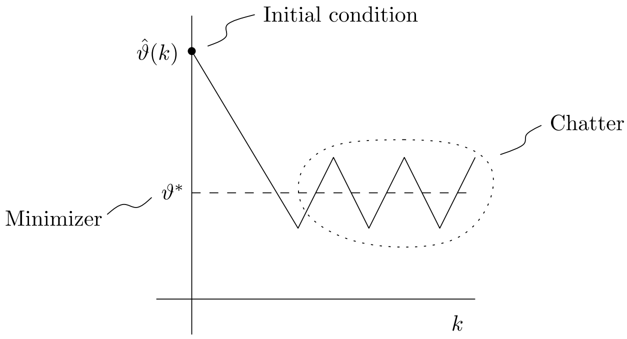

Figure 9. Illustration of chatter around an optimum.

4.4.1. Tuning of the gradient estimator. As discussed in section 4.2, the frequency of the perturbation $d(k)$ has to be sufficiently low such that the sawtooth period is close to steady-state at all time. In [18] it is shown that currently existing EC launchers, especially on smaller tokamaks like TEXTOR, operate on a longer time-scale than the sawtooth period dynamics. The speed of the launcher mirror is largely characterized by the bandwidth of its own motion control system. Roughly speaking, the launcher mirror can follow frequencies up to the bandwidth quite well (so that $\vartheta_a \approx \vartheta$) but attenuates higher frequencies (so that $\vartheta_a \approx 0$). The bandwidth of the launcher thus determines the maximal perturbation frequency.

This bandwidth f_{bw} is specified in the continuous-time domain, while the proposed ESC operates in discrete time. The fastest perturbation frequency in continuous time occurs if the sampling interval of the controller is the shortest, i.e. when the sawtooth period is the shortest. Hence, the perturbation frequency should be selected such that its continuous-time equivalent at the minimal sawtooth period is slightly lower than the bandwidth of the launcher. Defining $\tau_{s,min}$ as the smallest possible sawtooth period, the perturbation frequency a should thus satisfy

$$a \leqslant 2 f_{bw} \cdot \tau_{s,min}. \tag{13}$$

The earlier posed limits on parameter a still apply as well, i.e. $0 < a < 1$ and $2/a$ must be a natural number. With a realistic bandwidth f_{bw} of 10 Hz [18] and a minimal sawtooth period of 3.0 ms the perturbation frequency would be $a = 0.05\ \text{crash}^{-1}$. This perturbation frequency ensures optimal performance of the gradient estimator, assuming that the launcher dominates the sawtooth period dynamics.

In case the EC launcher operates on a faster time-scale than the sawtooth dynamics, the selection of parameter a requires identification of the slowest time-scale of the sawtooth period. In that case the perturbation should belong to the pass-band of the sawtooth period dynamics. Stepwise identification experiments as described in [18] give information on the time-scale of the sawtooth period. The period of the perturbation $d(k)$ should be larger than the settling time of such a step. In [18], a step around $\vartheta = 2.68°$ has a settling time of five crashes. Taking a safety margin into account, a perturbation period of ten crashes could be chosen. Simulations indeed showed that, in the case of an infinitely fast launcher, the perturbation frequency could be increased to $a = 0.2\ \text{crash}^{-1}$, corresponding to a ten crashes period.

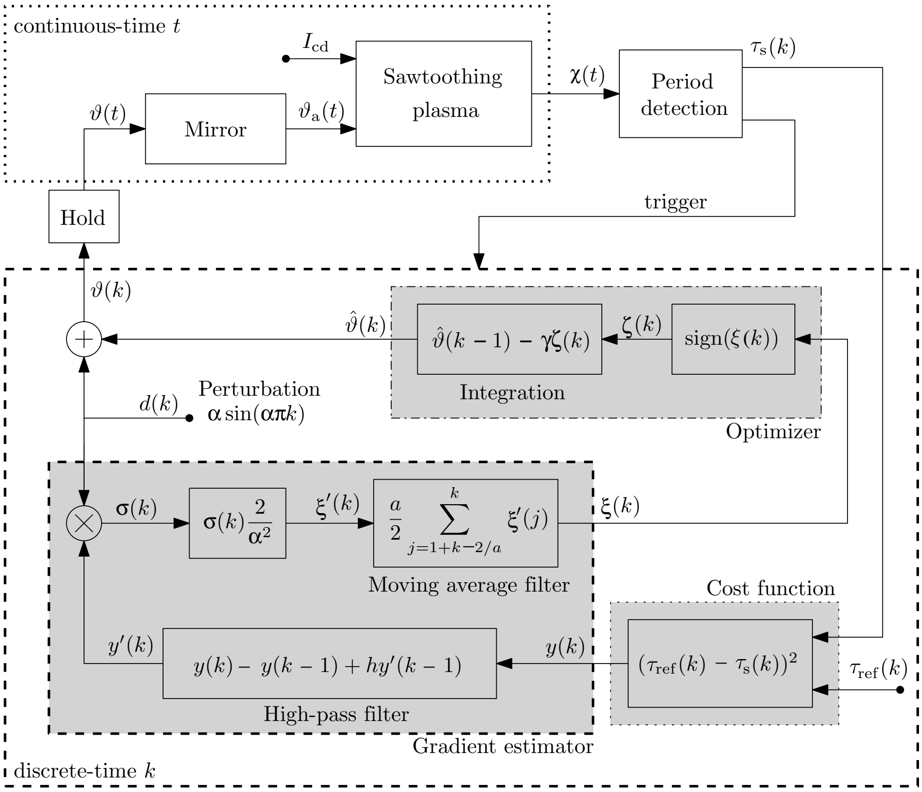

Figure 10. Complete closed loop with the ESC subsystems, i.e. the cost function, the gradient estimator, and the optimizer, marked grey.

The amplitude α of the perturbation $d(k)$ must be chosen as small as possible, since larger α lead to larger estimation errors in equation (8). However, the minimal amplitude is often limited by practical considerations. When disturbances are large, the gradient estimator can become quite slow, as it then takes more time to average out the effect of these disturbances. A larger perturbation amplitude can then improve the signal-to-noise ratio, and hence the speed, of the gradient estimation $\xi(k)$. Similarly, whenever hysteresis is present [15], α should be large enough to identify the frequency of the perturbation $d(k)$ in the time-evolution of the sawtooth period; this way the sign of the gradient can still be obtained, despite the hysteresis. In a practical setup the positioning accuracy of the mirror is likely to pose a lower limit on the minimal perturbation amplitude. E.g. the mirror on TEXTOR has a maximal positioning error of 0.6° during maximal acceleration [32], and smaller positioning errors during normal motion. Here the perturbation amplitude is chosen $\alpha = 0.3°$.

The high-pass filter in figure 10 is of a discrete-time first-order type. It has a single tuning parameter h which must satisfy $0 < h < 1$. The lower the value of h the more aggressive the filtering, i.e. for very low h the output of the filter would always be close to zero and a lot of the content of $y(k)$ is lost. It is therefore important to choose h close to 1. A practical value is $h = 0.9$.

4.4.2. Tuning of the optimizer. The optimizer gain γ is directly related to the optimization speed and should be tuned such that the optimizer is slower than the gradient estimator. If the gain is chosen too large, the optimizer tends to overshoot and miss the optimum since the gradient estimator is not able to determine the gradient accurately if the operating point is changing too fast. A pragmatic approach is to start with a small γ and increase the gain during experiments. Here it is assumed that the gradient estimator needs three perturbation periods to settle, i.e. $3 \cdot 2/a$ sawtooth crashes to obtain a proper gradient estimate, since the moving average filter has a memory of one perturbation period and the high-pass filter introduces some additional settling time. The estimated minimizer $\hat{\vartheta}(k)$ may not change too much in this time window. Figure 2 shows that the largest gradient is around $\vartheta \approx 7.5°$, a region which is approximately 2° wide. We therefore assume that the mirror angle may vary about 0.5° during the settling of the gradient estimator. An initial guess for the optimizer gain is thus

$$\gamma = \frac{0.5}{6} a \quad [°/\text{crash}]. \tag{14}$$

With $a = 0.2$ crash^{-1} (assuming an infinitely fast launcher) the initial guess is $\gamma = 0.017°$/crash, in simulations we could choose $\gamma = 0.02°$/crash.

Note that the total closed-loop error comprises the sum of the effects of both the perturbation $d(k)$ and the chatter. When operating close to the marginal point for NTM seeding, one should realize that this error could lead to an accidentally long, NTM triggering, sawtooth period. In such situations the perturbation and chatter amplitudes should be tuned with care.

4.5. Usage of a supervisor for practical issues

In most control implementations supervisory controllers are employed to detect and correct undesired closed-loop

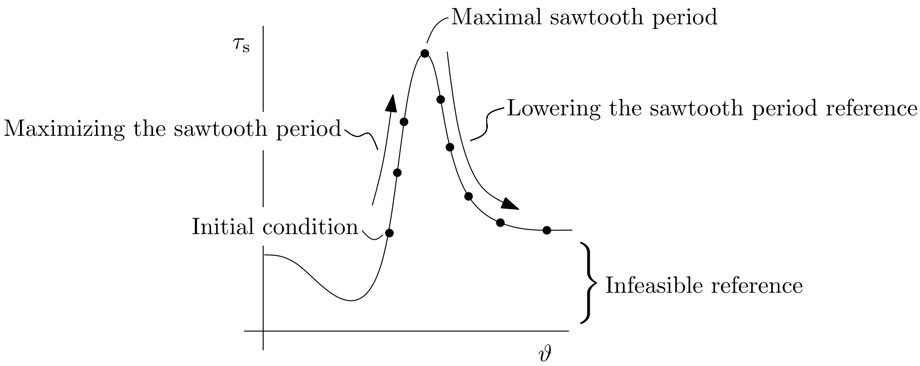

Figure 11. Unbounded mirror angle growth due to an infeasible reference.

behaviour. For the sawtooth control problem a supervisor can be used to guard the low-level controller, i.e. the ESC, since the requested mirror angle could grow unbounded for certain specific reference trajectories.

This is illustrated in figure 11. If the sawtooth period is first maximized and then sequentially lowered, it is not *a priori* known in which direction the ESC will choose to evolve, i.e. either to the left or to the right of the maximum. If the ESC steers towards the right it is possible to have unbounded growth of the mirror angle if $\tau_{s,ref}$ is set to a value smaller than the ohmic sawtooth period. A supervisor can be used to detect such unbounded growth and steer the mirror back using feedforward techniques.

Detection could be done by means of an additional gradient estimator, which estimates the gradient of the input–output map of figure 2 directly (instead of the gradient of the cost function). If the sign of this gradient is negative for mirror angles beyond the $q = 1$ surface, unbounded growth might occur, and the mirror needs to be steered back accordingly. Another solution is to simply limit the maximally allowed mirror angle. Similarly, a supervisor can prevent convergence to a non-optimal local minimum at $\vartheta = 0°$, e.g. by steering the angle to an arbitrary value outside $q = 1$ whenever $\tau_{s,ref} > \tau_\Omega$. The actual design of a supervisory controller is out of scope for this paper, as all simulation results are obtained without the use of a supervisor.

5. Simulation results

To validate the effectiveness of the proposed ESC strategy, this section shows and discusses the following closed-loop simulations:

(1) the tracking of sawtooth period reference trajectories in section 5.1;
(2) testing the controller for robustness against varying plasma parameters and disturbances in section 5.2;
(3) the impact of the launcher dynamics on the performance in section 5.3.

For cases 1 and 2 the response of the launcher mirror to a requested angle is assumed to be infinitely fast compared with the sawtooth dynamics, which is a reasonable assumption for very large tokamaks such as ITER. The controller parameters are the same for all simulations in section 5.1 and section 5.2, in section 5.3 the tuning is adjusted to cope with the launcher dynamics.

5.1. Tracking of sawtooth period references

In the following results the EC driven current $I_{CD} = 2$ kA, the perturbation frequency $a = 0.2$ crash^{-1}, its amplitude $\alpha = 0.3°$, the optimizer gain $\gamma = 0.02°$/crash and the high-pass filter parameter $h = 0.9$, as has been discussed in section 4.4.

5.1.1. Tracking small sawtooth periods. In the first simulation the initial mirror angle is chosen $\vartheta(k = 0) = 3°$ and the reference trajectory for the sawtooth period takes values between 5 and 14 ms, i.e. the controller operates on the left-hand side of figure 2. The closed-loop results are shown in figure 12. The sawtooth period as a function of the elapsed sawtooth crashes is shown in figure 12(*a*), the reference is indicated with the thicker grey line. Figure 12(*b*) shows the mirror angle and the estimate of the minimizer $\hat{\vartheta}$, which shows the chatter due to the sliding mode optimizer. In figure 12(*c*), the error $\tau_s - \tau_{s,ref}$ is shown, which is the difference between the actual sawtooth period and the reference. The gradient estimate ξ is shown in figure 12(*d*). The reference starts at 10 ms, and is then gradually lowered to 5 ms. Then the reference sawtooth period is increased, first linearly then in a stepwise fashion. The results show that there is a small transient from the initial condition. The controller achieves tracking of the requested sawtooth period with an error of slightly more than 1 ms when the reference is changing gradually. The controller successfully handles a step in the reference applied at $k = 800$ crashes; the settling takes place in about 100 crashes.

The oscillations on the sawtooth period are a result of the perturbation on ϑ and the chatter on $\hat{\vartheta}$ introduced by the sliding mode optimizer. Since the gradient of the input–output map is relatively small, the resulting oscillations on τ_s are small. In figure 12(*b*) the amplitude of the chatter is approximately equal to the perturbation $d(k)$.

The step in $\tau_{s,ref}$ leads to a large increase in the gradient estimation in figure 12(*d*). A first-order gradient descent optimizer equation (4) as in [20], whose optimization rate is proportional to the gradient, would have yielded a significant change in the mirror angle at this step. Such a change in ϑ can be risky, as it can instantly bring the sawtooth system to another operating region (e.g. the right-hand side of figure 2), which can result in an unbounded growth of ϑ. In a specific incidental case, such a change in ϑ may also turn out just right, yielding

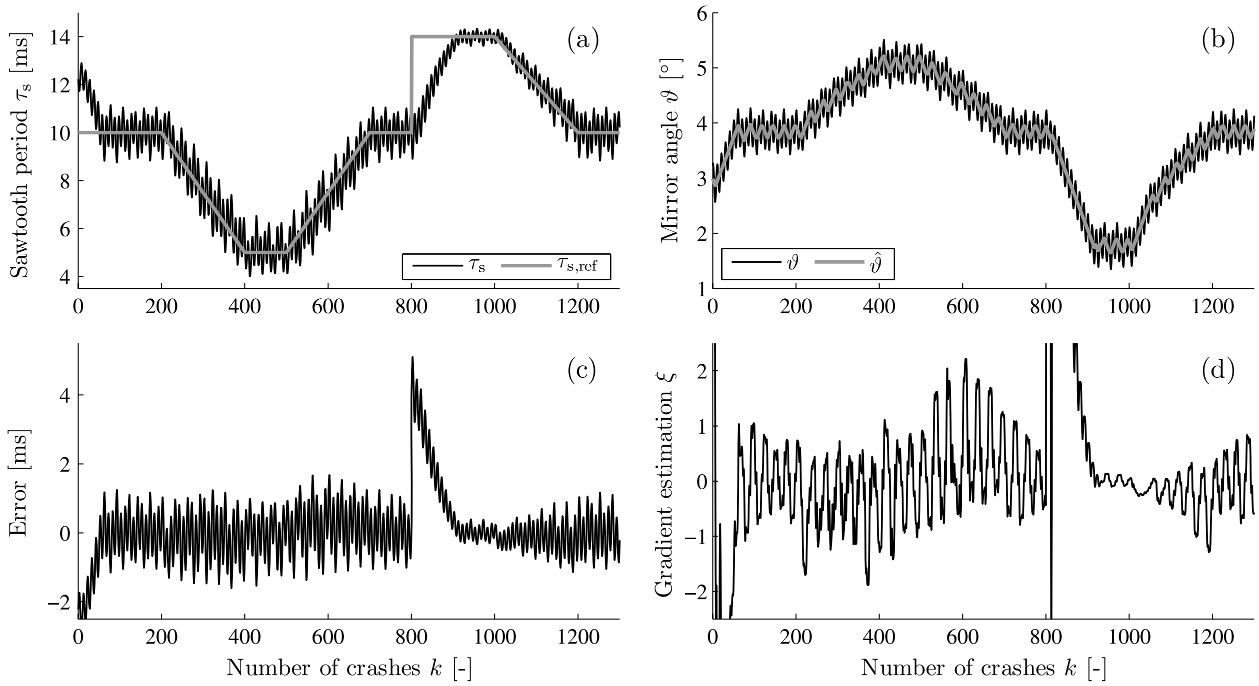

Figure 12. Simulation results of tracking small sawtooth period references: (*a*) the sawtooth period τ_s, (*b*) the mirror angle ϑ and estimated minimizer $\hat{\vartheta}$ which shows the chatter, (*c*) the tracking error $\tau_s - \tau_{s,ref}$ and (*d*) the gradient estimate.

an uncharacteristically fast convergence for ESC. An example can be found in [20], where the controller adapts within just one perturbation period to a rapid change in vertical position. In contrast, figure 12(*b*) shows that the sliding mode optimizer always alters ϑ linearly towards the new optimum. Hence, due to this constant optimization rate the sliding mode optimizer is more robust against large variations of the gradient.

5.1.2. Tracking large sawtooth periods. Figure 13 shows the simulation result where the initial mirror angle $\vartheta(k = 0) = 7.3^\circ$ and the reference takes values between 8 and 25 ms, including a step at $k = 800$ crashes. The controller operates in the middle region of figure 2, typically around a mirror angle of 7.5°. There is a small transient from the initial condition, after which the controller achieves good tracking. The step in the reference is handled successfully as well. Although the step is twice as large as the previous result, the settling time is only 30 crashes. This is because the gradient in this operating region is much larger than compared with the smaller sawteeth in figure 12. Hence, small changes in ϑ have a large effect on τ_s. So although the convergence rate in °/crash is fixed, τ_s converges much faster, since a much smaller change in mirror angle is required. However, the large gradient has also caused the oscillations on the sawtooth period to increase by a factor 5. To reduce these oscillations one could make the amplitude of the perturbation smaller in regions where the estimated gradient is large, e.g. by means of gain-scheduling as suggested in [20].

5.1.3. Tracking in different operating regions. One of the advantages of the proposed ESC is that it can switch between regions with different signs of the input–output map gradient, so that one can cover the whole operating region with a single controller. The simulation results in figure 14 demonstrate this. Starting from the right-hand side of the input–output map where $\vartheta(k = 0) = 10^\circ$, the ESC first maximizes the sawtooth period to the steady-state maximum of about 30 ms, as the 31 ms reference value is unreachable. The controller then quickly decreases the mirror angle to reach the 10 ms reference value and tracks successfully as the reference gradually changes to 5 ms. Figure 14(*a*) shows that the period makes a small jump around $k = 600$ crashes; there the controller jumps from the middle region (large gradient) to the left-hand-side region (smaller gradient). This can be explained by figure 14(*c*), which depicts the sawtooth period as a function of the mirror angle. The thickness of the band created by the trajectory of τ_s around the input–output map indicates the accuracy of the gradient estimate. For mirror angles $5.5^\circ \leqslant \vartheta \leqslant 6.5^\circ$ the band is very narrow, which implies that the sawtooth dynamics is very fast in this region (as has been shown in [18]). The gradient estimation is therefore very accurate, which makes this region attractive. A wider band, as for $\vartheta > 6.5^\circ$, indicates a less accurate gradient estimation due to slow sawtooth dynamics. The trajectory of τ_s at $\vartheta \approx 6.7^\circ$ is actually perpendicular to the steady-state input–output map, thanks to which the controller crosses the minimum to the region around $\vartheta = 6^\circ$. Hence, the ESC converges to the smallest of the two angles corresponding to 5 ms, which is actually quite eligible since the oscillations on τ_s are much smaller there. Afterwards, the ESC continues to decrease the mirror angle to eventually reach 10 ms again.

In our case the region with slow dynamics is surrounded by fast sawtooth dynamics. Hence, the somewhat erroneous gradient estimations in the slow region are intercepted by the neighbouring fast dynamics, where the gradient is estimated correctly. It should be noted that the slow response around $\vartheta \approx 6.7^\circ$ and the fast one around $\vartheta \approx 6^\circ$ are predictions of the model in section 2.1, and may or may not occur in real

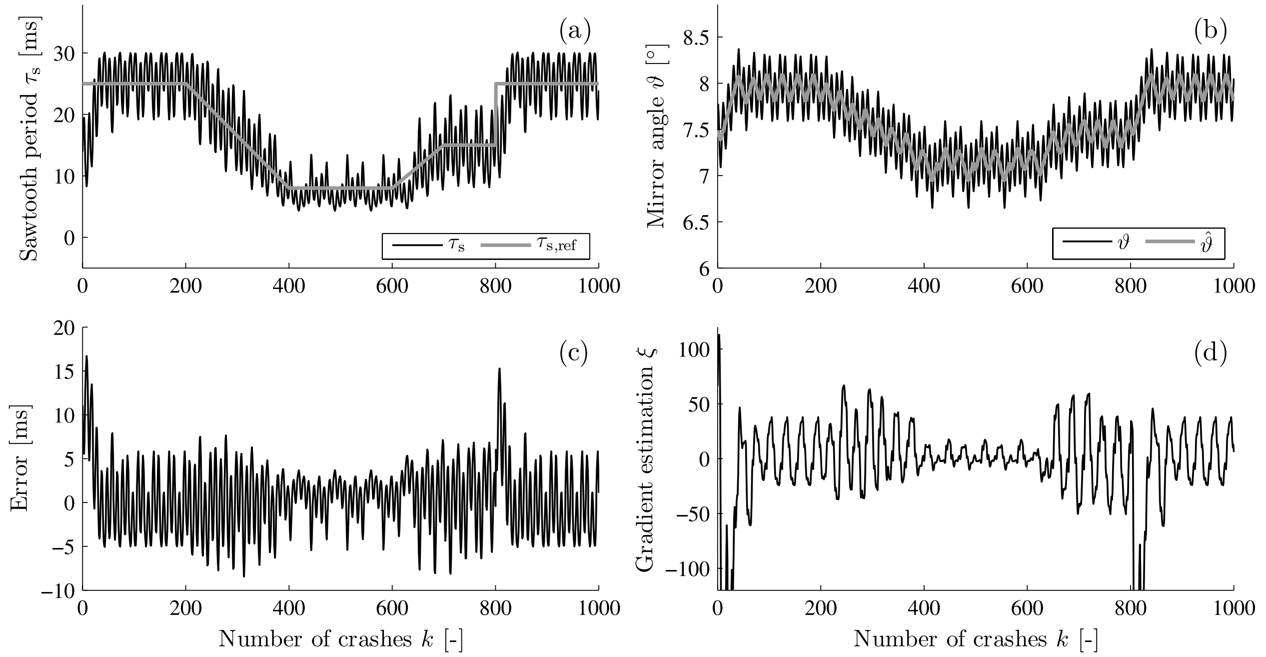

Figure 13. Simulation results of tracking large sawteeth: (*a*) the sawtooth period τ_s, (*b*) the mirror angle ϑ and estimated minimizer $\hat{\vartheta}$ with chatter present, (*c*) the tracking error $\tau_s - \tau_{s,ref}$ and (*d*) the gradient estimate ξ.

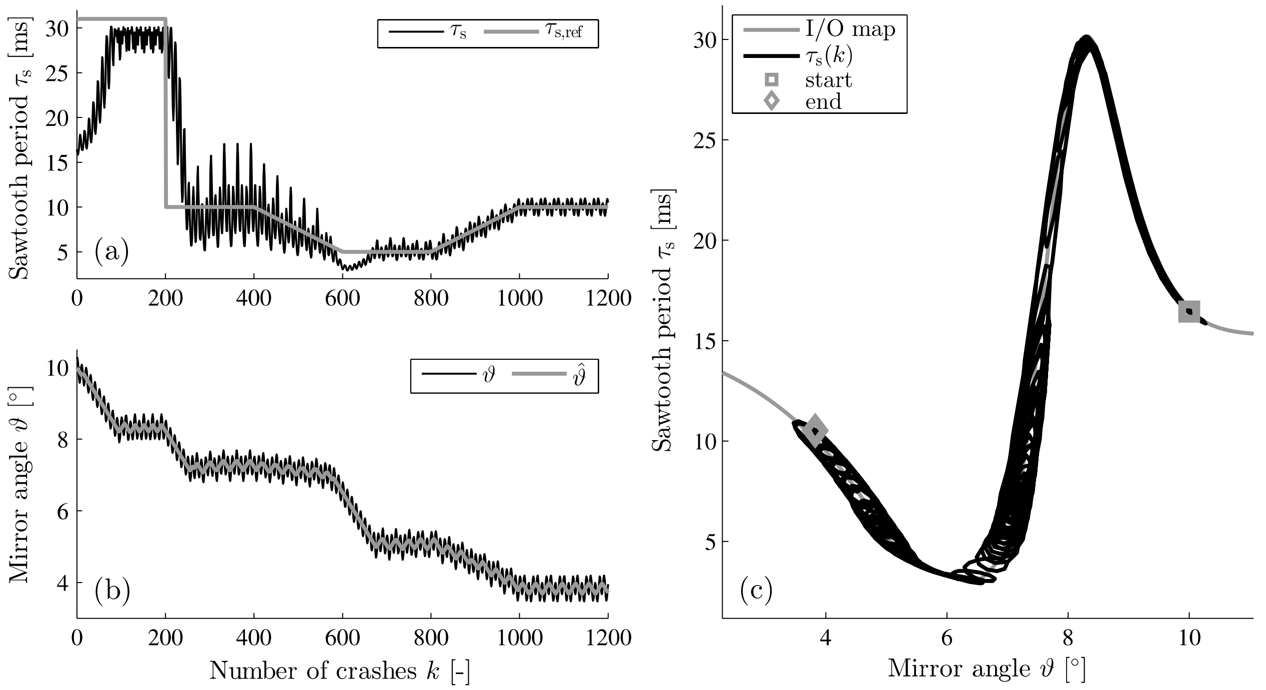

Figure 14. Simulation results of sawtooth period tracking in three different regions: (*a*) the sawtooth period τ_s, (*b*) the mirror angle ϑ and estimated minimizer $\hat{\vartheta}$ and (*c*) the input–output map and sawtooth period as function of the mirror angle. The markers indicate the start and the end of the sawtooth period trajectory.

experiments. In practice, one may encounter other phenomena that can have effects on the behaviour of the controller. The simulations show that ESC can handle such phenomena, as long as the separation of time-scales is satisfied.

5.2. *Extremum seeking and robustness*

For experimental devices controller robustness is an important issue. Robustness is considered as the ability of the controller to handle disturbances and uncertainties. The simulations in this section will indeed demonstrate the robustness of ESC; firstly, by considering variations of plasma parameters, and secondly, by disturbing the closed loop with noise and delay.

5.2.1. Robustness against varying plasma parameters. First, the robustness of ESC is demonstrated by applying the same controller to a sawtooth model with different parameters; the driven current I_{CD} is changed from 2 to 1.8 kA, the plasma

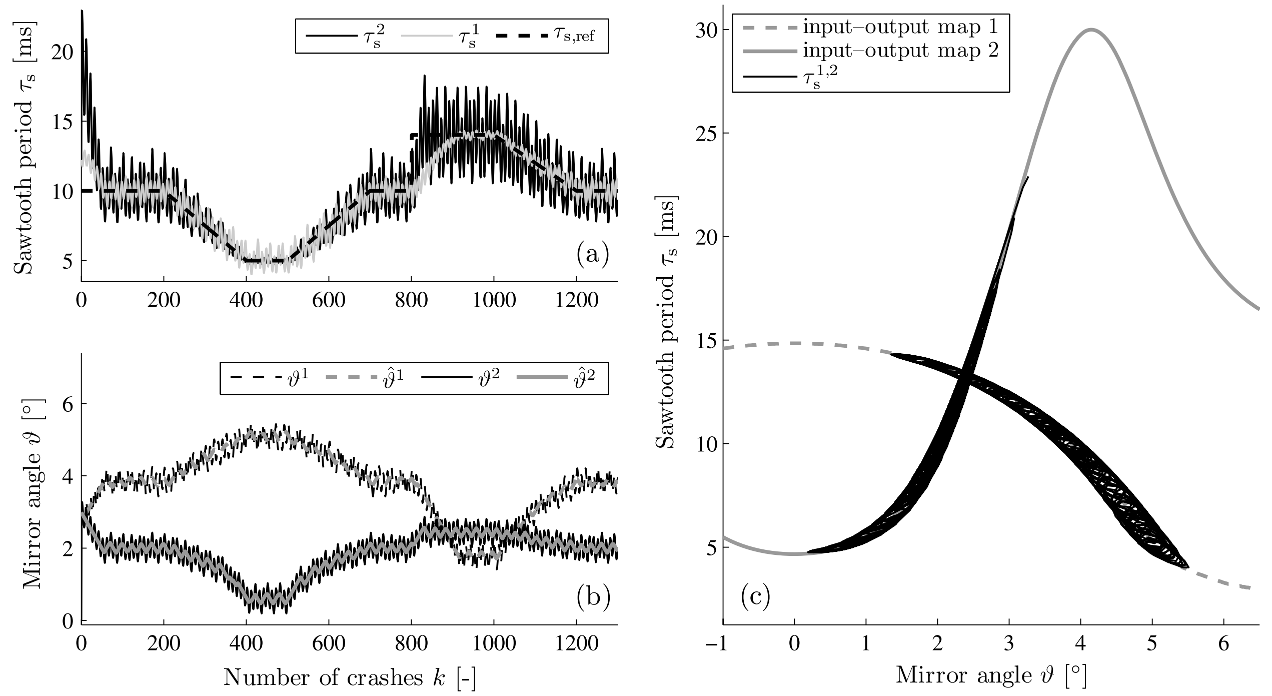

Figure 15. Simulation results with different plasma parameters ([2]), together with the original simulation indicated with ([1]). The sawtooth period τ_s is shown in (*a*), the mirror angle ϑ and estimated minimizer $\hat{\vartheta}$ in (*b*) and the sawtooth period trajectory in (*c*). We see that in spite of a dramatic variation of the gradients, including a change of sign, the sawtooth controller still performs robustly.

current I_p from 400 to 350 kA and the toroidal magnetic field B_ϕ from 2.45 to 2.4 T. The simulation results are shown in figure 15, together with the previous results from figure 12. Figure 15(*c*) shows that the steady-state input–output map has completely changed, due to the large shift of the $q = 1$ surface. Nevertheless, the reference is still successfully tracked. The oscillation on the sawtooth period at $\tau_s \approx 15$ ms is larger than before, since the gradient at this operating point is now larger. At $\tau_s \approx 5$ ms the oscillation is smaller, since that specific operating point now has a nearly zero gradient. The settling time after the step at 800 crashes is now much faster, which is again caused by the larger gradient at $\tau_s \approx 15$ ms.

A linear controller as used in [15, 16, 18] is not able to cope with such large parameter variations, as they change the sign of the gradient over a large portion of the operating space. ESC automatically adapts to such variations, even if these occur online, i.e. during a discharge. Such changes are similar to changes in the reference sawtooth period, since both introduce a sudden change in cost function. Hence, ESC is guaranteed to track these variations if they occur either on a slow enough time-scale, or stepwise with sufficient time between the steps (as in previous simulations). Note that if the parameter variations largely affect the sawtooth period dynamics, the time-scales of the controller subsystems might have to be adjusted.

5.2.2. Robustness against disturbances and detection delay. Real experiments are often subject to actuator and sensor noise. In the sawtooth control loop there can be noise on the amount of EC driven current and disturbances on the deposition location, e.g. due to EC beam deflections or fluctuations of the plasma position and shape. On the sensor side a sawtooth crash detection is never flawless and will inevitably introduce delay and noise on the sawtooth period. To demonstrate that ESC can cope with these disturbances the following disturbances have been added in a repeat of the simulation depicted in figure 12:

- a $\pm 0.25°$ uniform random noise on ϑ;
- a ± 20 A uniform random noise on I_{CD} (nominal value is 2 kA);
- a ± 0.5 ms uniform random noise and an additional delay of 5 ms on τ_s.

The results depicted in figure 16 show that the controller still successfully tracks the reference and is thus robust for the added disturbances. In principle the 5 ms detection delay introduces a phase shift between the perturbation $d(k)$ and the resulting sawtooth period output $\tau_s(k)$, which influences the gradient estimate. However, here this phase shift is at most one crash, which is quite small compared with the perturbation period of ten crashes. Moreover, the other disturbances on the gradient estimate are mostly attenuated by the moving average filter.

5.3. Impact of slow launcher dynamics

In the previous simulations the launcher mirror has been considered infinitely fast. In practice, this launcher is a closed-loop controlled mechanical device with a limited bandwidth, introducing dynamics between ϑ and the actual mirror angle ϑ_a. This dynamics can play an important role, especially when the launcher time-scale is longer than the sawtooth period. In this section the launcher closed loop is modelled in the same fashion as in [18], i.e. by means of a second order low-pass filter with a cut-off frequency (bandwidth) of $f_{\mathrm{bw}} = 10$ Hz and a damping of 0.35. With this model two simulations have

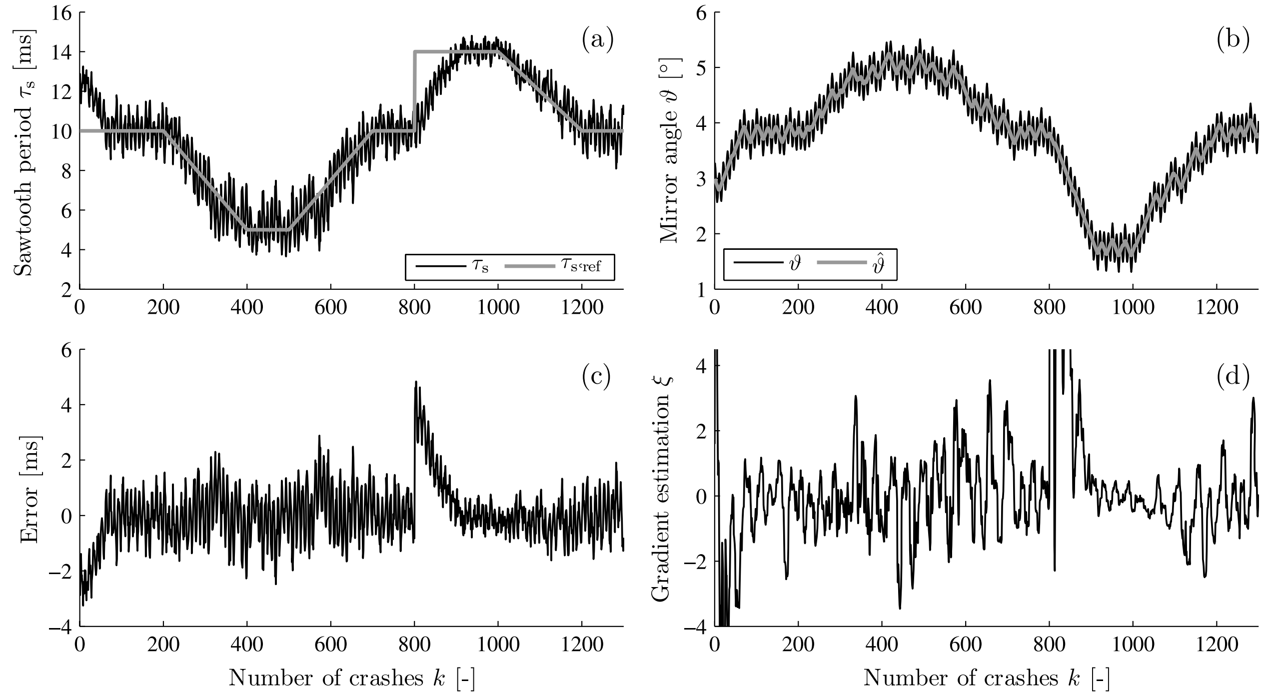

Figure 16. Simulation results with disturbances present: (*a*) the sawtooth period τ_s, (*b*) the mirror angle ϑ and estimated minimizer $\hat{\vartheta}$, (*c*) the tracking error $\tau_s - \tau_{s,\mathrm{ref}}$ and (*d*) the gradient estimate ξ.

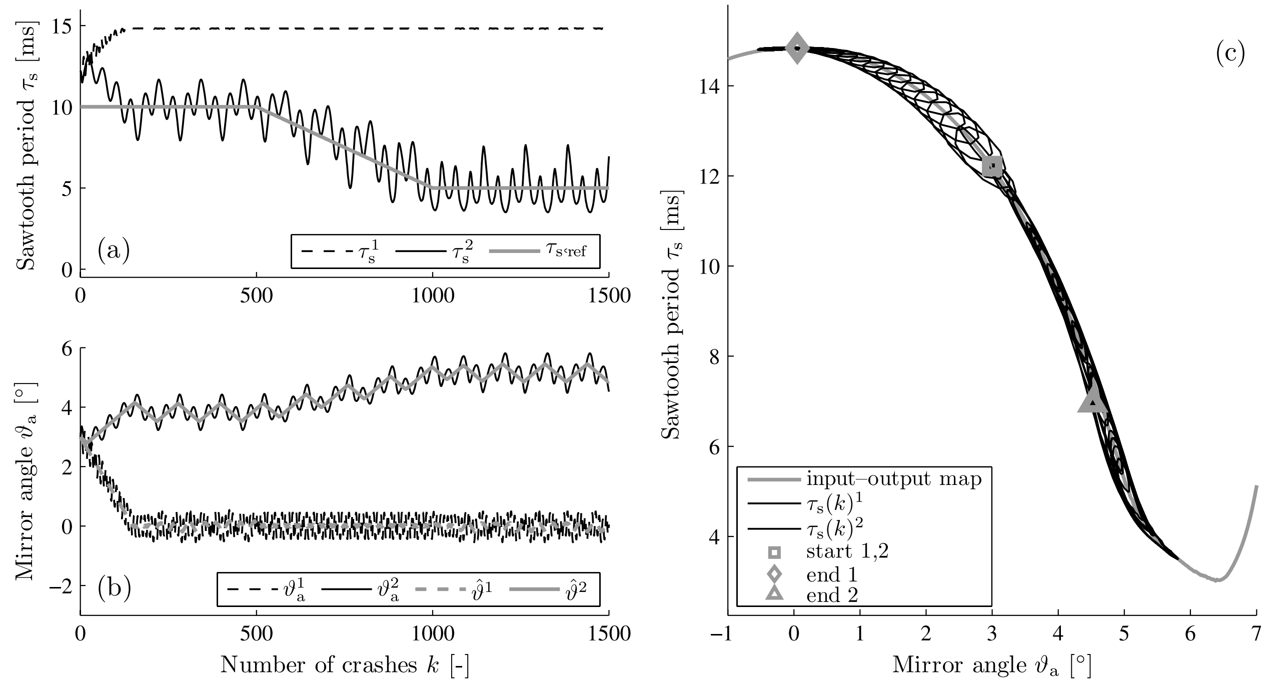

Figure 17. Simulation results with slow launcher dynamics, (^{1}) is with original tuning, (^{2}) is with adjusted tuning: (*a*) the sawtooth period τ_s, (*b*) is the actual mirror angle ϑ_a and estimate of the minimizer $\hat{\vartheta}$, and (*c*) is the trajectory of τ_s on the steady-state input–output map. Markers indicate the start and end of these trajectories.

been performed to investigate the influence of a slow launcher: one with the original controller tuning, and one with adjusted tuning. The results are shown in figure 17. The reference starts at 10 ms and is lowered to 5 ms. The original tuning, indicated with (^{1}), fails as it maximizes the cost function instead, due to a wrong estimation of the sign of the gradient. In the second simulation, indicated with (^{2}), the perturbation frequency is lowered to the bandwidth of the launcher (i.e. $a = 0.05$ crash^{-1}), as described in section 4.2, and the optimizer gain γ is reduced to $\gamma = 0.01°$/crash to preserve the separation of time-scales. With this adjusted tuning the functionality of the controller is restored. Note that due to the slow launcher the overall control loop is now a factor two slower.

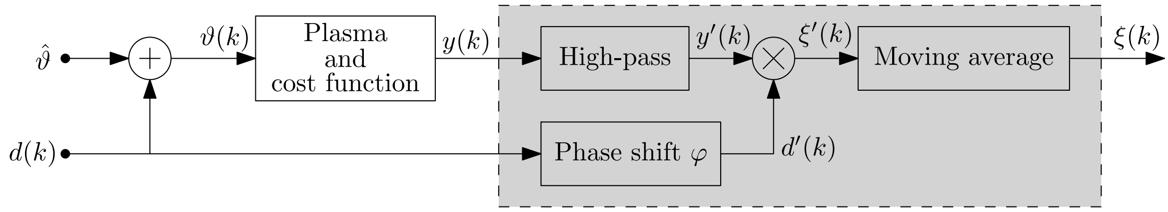

Figure 18. Gradient estimator topology with additional phase lag compensation.

With the adjusted tuning the settling time of the gradient estimator has increased. This is visible in the chatter, which is of lower frequency. The gradient estimator is four times as slow as before, the optimizer gain has been reduced by a factor two. Consequently, the amplitude of the chatter has increased with a factor two.

6. Extremum seeking and performance

The above results clearly demonstrate the extremely high robustness of the extremum seeking controller. This robustness is inherent to the fact that ESC is completely model-free; it simply estimates the momentary behaviour of the sawtooth period online. An ESC approach is therefore very useful under experimental tokamak conditions, where the exact plasma parameters are somewhat uncertain or conditions may vary to a large extent. The disadvantage of ESC is its relatively low performance. The simulations have shown that convergence (or settling time) can take more than 100 crashes (i.e. more than 1 s in continuous time), and the sawtooth period will always oscillate around its desired value due to the applied perturbation $d(k)$ and the chatter. Consequently, steady-state errors are about ± 1 ms for small τ_s, but can increase up to ± 5 ms for larger τ_s. For comparison, linear feedback controllers as in [18] and high-performance approaches as in [19] obtain zero steady-state errors, and the latter typically converges within 0.2 to 0.35 s. Hence, under predictable reactor-like conditions such approaches might be more beneficial.

The above illustrates the classical trade-off between robustness and performance. It also indicates that the performance of ESC can be improved by including explicit knowledge of the sawtooth into the controller, at the expense of some robustness. In this section we briefly suggest three methods to achieve this.

6.1. Feedforward control

In situations where the steady-state input–output map of the sawtooth period is approximately constant and *a priori* known, this knowledge can be employed in advance to steer the mirror angle towards a value corresponding to a desired $\tau_{s,\mathrm{ref}}$. This open-loop technique is called feedforward control, a widely used technique to improve closed-loop performance, such as in [19]. Feedforward control can easily be combined with ESC using the input–output map of figure 2 as a look-up table for the mirror angle and adding this to the value of ϑ computed by the optimizer. This will improve the settling time to a large extent (up to the time-scale of the launcher mirror), but it does not affect the steady-state oscillations. The extremum seeker will still perturb and chatter to make sure the system remains at the optimum of the cost function.

Feedforward control yields the best results when the input–output map is known accurately. If the calculated feedforward angle is significantly wrong, it can be viewed as a disturbance or change in cost function. Since ESC is robust against such disturbances, it will compensate for the error in the feedforward, and still steer towards the desired sawtooth period. However, this does reduce the performance in terms of the settling time, possibly making it worse than without feedforward.

6.2. Phase lag compensation

The sinusoidal perturbation on ϑ leads to an oscillation of the sawtooth period τ_s with the same base frequency as ϑ. For slow perturbations these signals are approximately in phase, but for higher frequencies the sawtooth (and launcher) dynamics can introduce a phase shift between ϑ and τ_s, which influences the accuracy of the gradient estimate. In case of a phase shift ψ it can be shown that the averaged gradient estimate becomes

$$\lim_{K\to\infty} \frac{1}{K}\sum_{k=0}^{K}(\xi(k)) = \cos(\psi)\cdot\frac{\alpha^2}{2}\,\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\bigg|_{\vartheta=\hat{\vartheta}}. \qquad (15)$$

Hence, a mismatch $\cos(\psi)$ is introduced, which can lead to a wrong estimation of the sign of the gradient when $\psi \geqslant \pi/2$. This can be compensated for by phase-shifting the demodulation signal with an estimate of the phase lag [33] as shown in figure 18. Let φ be the phase shift in the demodulation signal. The gradient estimate then becomes

$$\lim_{K\to\infty} \frac{1}{K}\sum_{k=0}^{K}(\xi(k)) = \cos(\psi-\varphi)\cdot\frac{\alpha^2}{2}\,\frac{\mathrm{d}f}{\mathrm{d}\vartheta}\bigg|_{\vartheta=\hat{\vartheta}}.$$

Hence, if the demodulation phase shift φ equals the phase lag ψ, the gradient estimate is not affected. However, the phase lag ψ in the discrete-time domain, discussed in section 3, varies with the sawtooth period. In this case φ should represent the average phase lag, since $\varphi < \psi$ has the same effect on the gradient estimation as $\varphi > \psi$. Determining this average phase lag at the perturbation frequency requires dynamic identification experiments, as has been proposed in [18]. Alternatively, the demodulation signal $d'(k)$ can be constructed by 'holding' the perturbation $d(k)$, filtering it with a continuous-time model of the launcher dynamics and sampling the output at every sawtooth crash. This effectively compensates for the phase lag ψ.

When the phase lag is compensated for, the perturbation frequency can be increased, yielding shorter settling times of the gradient estimator. Consequently, the optimizer gain can be

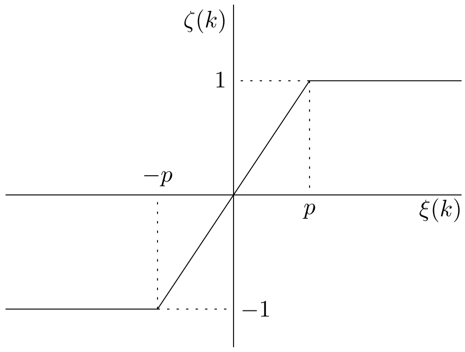

Figure 19. Saturation function and tuning parameter p.

increased, leading to faster convergence of the overall control loop.

6.3. Reducing the tracking error

Part of the oscillations on the error shown in the simulation results are introduced by chatter, which is a side effect of the chosen sliding mode optimizer. A common approach to eliminate chatter is the use of the saturation function (or a similar function like a hyperbolic tangent) instead of the sign function [34] in the optimizer equation (12). This way the optimization speed is constant when the estimated gradient is large, and proportional to the gradient when the gradient is small.

The saturation function depicted in figure 19 has limits $-1, 1$, as does the sign function, and has a parameter p that can be tuned to select at which value of the gradient to switch from sliding mode to first-order gradient descent. However, the simulation results in section 5.1 show that the gradient varies significantly with the operating point. If p is small compared with the gradient estimate, the optimizer would only operate like a sliding mode optimizer, and chatter would not be eliminated. If p is large compared with the gradient estimate, the optimizer only uses first-order gradient descent, which reduces robustness and affects the convergence speed. Hence, eliminating chatter improves the performance in terms of steady-state error, but at the expense of slower convergence and less robustness.

7. Conclusions

In this paper we have provided a structured design of a robust sawtooth period controller. The proposed extremum seeking controller is a special type of adaptive controller. Its working principle relies on online identification of the gradient of a cost function, which has a minimum at the desired sawtooth period, by a gradient estimator. The benefit of this control strategy is that it does not rely on any model describing the behaviour of the sawtooth period and is therefore highly robust. Moreover, in principle it can be implemented and tuned without any preceding identification experiments at all.

We have discussed each of the three building blocks of the controller, i.e. the cost function, the gradient estimator and the optimizer, and designed each one such that the sawtooth period is controlled with maximal robustness and acceptable performance. Furthermore, practical tuning guidelines for the controller parameters have been suggested, based on the required separation of time-scales between these subsystems, and several performance improvements have been introduced.

The behaviour of the extremum seeking controller has been assessed in a closed-loop interconnection with a Kadomtsev–Porcelli sawtooth model. Simulation results demonstrated the tracking of sawtooth period references in different operating regimes. The controller is able to handle stepwise changes in the sawtooth period reference, changes in plasma parameters, additional crash detection delay and disturbances on both the actuator and sensor side. In each case the controller tracks the desired sawtooth period, which demonstrates its high robustness.

ESC does come with a non-zero steady-state error, as the sawtooth period is always oscillating around the reference value. This is partly caused by an external sinusoidal perturbation on the EC mirror angle, which is needed for the online estimation of the gradient of the cost function. Moreover, there is an additional chatter on the mirror angle, introduced by the sliding mode optimizer.

The high robustness against plasma uncertainties and disturbances makes ESC applicable to a wide variety of sawtooth models and real sawtoothing plasmas, even when there is a discontinuity in the steady-state input–output map, which might arise in the presence of fast-ions [17]. ESC can easily go across such a discontinuity, which will be interpreted as a very large gradient, and stabilize the sawtooth period on either side of it (stabilization on the discontinuity itself will be subject to similar jumps as in [17]). ESC is therefore a particularly interesting candidate to control sawteeth on experimental devices, where plasma variations are large and performance requirements are low, but also during the commissioning phase of fusion reactors to search for suitable operating conditions. Under controlled plasma conditions where closed loop performance (in terms of convergence speed and steady-state error) is more important, such as in sawtoothing reactor scenarios, high-performance sawtooth control strategies as in [19] are probably more suitable. Finally, the possible applications of the proposed controller are not limited to the sawtooth control problem only; ESC has been used before in fusion research [35], and remains an interesting candidate for any problem in or around the tokamak with non-linearities, large model uncertainties or parameter variations.

Acknowledgments

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.

References

[1] von Goeler S., Stodiek W. and Sauthoff N. 1974 *Phys. Rev. Lett.* **33** 1201–3

[2] Hastie R.J. 1997 *Astrophys. Space Sci.* **256** 177–204
[3] Wesson J.A. and Campbell D.J. 2004 *Tokamaks* 3rd edn (Oxford: Oxford University Press)
[4] Nave M.F.F. *et al* 2003 *Nucl. Fusion* **43** 1204–13
[5] Sauter O. *et al* 2002 *Phys. Rev. Lett.* **88** 105001
[6] Gude A., Günther S., Maraschek M., Zohm H. and the ASDEX Upgrade team 2002 *Nucl. Fusion* **42** 833–40
[7] Buttery R., Hender T., Howell D., Haye R.L., Parris S., Sauter O., Windsor C. and JET-EFDA Contributors 2004 *Nucl. Fusion* **44** 678
[8] Chapman I. *et al* and the ASDEX Upgrade, DIII-D, HL-2A, JT-60U, MAST, NSTX, TCV and Tore Supra Teams and JET-EFDA Contributors 2010 *Nucl. Fusion* **50** 102001
[9] Porcelli F, Boucher D and Rosenbluth M N 1996 *Plasma Phys. Control. Fusion* **38** 2163–86
[10] Chapman I.T. 2011 *Plasma Phys. Control. Fusion* **53** 013001
[11] Mück A., Goodman T.P., Maraschek M., Pereverzev G., Ryter F., Zohm H. and Team A.U. 2005 *Plasma Phys. Control. Fusion* **47** 1633–55
[12] Merkulov A., Schüller F.C., Westerhof E., de Baar M.R., Krämer-Flecken A., Liang Y. and TEXTOR Team 2004 *Proc. of Joint Varenna–Lausanne Int. Workshop on Theory of Fusion Plasmas (Varenna, Italy, 30 August–3 September 2004)* p 279
[13] Angioni C., Goodman T.P., Henderson M.A. and Sauter O. 2003 *Nucl. Fusion* **43** 455–65
[14] van Berkel M., Witvoet G., de Baar M.R., Nuij P.W.J.M., ter Morsche H.G. and Steinbuch M. 2011 *Fusion Eng. Des.* **86** 2908–19
[15] Paley J.I., Felici F., Coda S., Goodman T.P., Piras F. and the TCV Team 2009 *Plasma Phys. Control. Fusion* **51** 055010
[16] Lennholm M. *et al* 2009 *Fusion Sci. Technol.* **55** 45–55
[17] Lennholm M. *et al* 2009 *Phys. Rev. Lett.* **102** 115004
[18] Witvoet G., de Baar M.R., Westerhof E., Steinbuch M. and Doelman N.J. 2011 *Nucl. Fusion* **51** 073024
[19] Witvoet G., Steinbuch M., de Baar M.R., Doelman N.J. and Westerhof E. 2012 *Nucl. Fusion* **52** 074005
[20] Paley J.I., Felici F., Coda S., Goodman T.P. and the TCV Team 2009 *Plasma Phys. Control. Fusion* **51** 124041
[21] Tan Y., Moase W.H., Manzie C., Nešić D. and Mareels I.M.Y. 2010 *Proc. 29th Chinese Control. Conf. (Beijing, China, 29–31 July 2010)* pp 14–26 http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5572972
[22] Krstić M. and Wang H.H. 2000 *Automatica* **36** 595–601
[23] Krstić M. 2000 *Syst. Control Lett.* **39** 313–26
[24] Nešić D. 2009 *Eur. J. Control* **15** 331–47
[25] Nešić D., Tan Y., Moase H. and Manzie C. 2010 *Proc. 49th Conf. Decision Control. (Atlanta, GA, 15–17 December 2010)* pp 4625–30 http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5717929
[26] Pan Y., Özgüner U. and Acarman T. 2003 *Int. J. Control.* **76** 968–85
[27] Kadomtsev B.B. 1975 *Sov. J. Plasma Phys.* **1** 389
[28] Neubauer O., Czymek G., Giesen B., Hüttemann P.W., Sauer M., Schalt W. and Schruff J. 2005 *Fusion Sci. Technol.* **47** 76–86
[29] Tan Y., Nešić D. and Mareels I. 2006 *Automatica* **42** 889–903
[30] Choi J.Y., Krstić M., Ariyur K.B. and Lee J.S. 2002 *IEEE Trans. Autom. Control.* **47** 318–23
[31] Tan Y., Nešić D. and Mareels I. 2008 *Automatica* **44** 1446–50
[32] Hennen B.A., Westerhof E., Oosterbeek J.W., Nuij P.W.J.M., Lazzari D.D., Spakman G.W., de Baar M., Steinbuch M. and the TEXTOR team 2009 *Fusion Eng. Des.* **84** 928–34
[33] Elong E., Krstić M. and Ariyur K.B. 2000 *Proc. Am. Control. Conf.* (*Chicago, IL, 28–30 June 2000)* pp 428–32 http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=878936
[34] Slotine J.J.E. and Li W. 1991 *Applied Nonlinear Control* (Englewood Cliffs, NJ: Prentice Hall)
[35] Felici F. *et al* and the LHD Experiment Group 2010 *Nucl. Fusion* **50** 105003