

A dither-free extremum-seeking control approach using 1st-order least-squares fits for gradient estimation*

B.G.B. Hunnekens¹, M.A.M. Haring², N. van de Wouw¹ and H. Nijmeijer¹

Abstract—In this paper, we present a novel type of extremum-seeking controller, which continuously uses past data of the performance map to estimate the gradient of this performance map by means of a 1st-order least squares fit. The approach is intuitive by nature and avoids the need of dither in the extremum-seeking loop. The avoidance of dither allows for an asymptotic stability result (opposed to practical stability in dither-based schemes) and, hence, for exact convergence to the performance optimal parameter. Additionally, the absence of dither eliminates one of the time-scales of classical extremum-seeking schemes, allowing for a possibly faster convergence. A stability proof is presented for the static-map setting which relies on a Lyapunov-Razumikhin type of proof for time-delay systems. Simulations illustrate the effectiveness of the approach also for the dynamic setting.

I. INTRODUCTION

Extremum-seeking control is an adaptive control approach that optimizes a certain performance measure in terms of the steady-state output of a system in real-time, by automated adaptation of the system parameters [1], [13], [12], [6]. Different classes of extremum-seeking approaches exist in the literature: the classical approaches which continuously adapt the system parameters in order to estimate the low-order derivatives (often only the gradient) of the performance map [13], [12], [6], [10], and numerical optimization-based methods [9], [14], [4]. In this paper, we will focus on a continuous-adaptation-based approach.

Commonly, in these continuous-adaptation-based extremum-seeking works, an external dither-signal is injected into the extremum-seeking loop to enable the estimation of the derivatives of the performance map. Classical schemes use modulation of the dither signal and the measured performance signal in order to obtain an estimate of the derivatives [13], [12], [6], [10]. Although such an approach for estimating the derivatives works well under suitable conditions, and has been applied successfully in many practical applications, there are also some drawbacks. Firstly, the application of the dither generally hampers the true convergence to the performance-optimal setting, characterized by *practical* stability results in literature (see e.g. [6], [13]) in contrast to desired *asymptotic* stability results. Secondly, the dither signal constitutes one

of the important time-scales in the extremum-seeking scheme. Since this time-scale should be separated from the time-scale of the plant-dynamics and the optimizer, this may limit the convergence speed of the algorithm.

In this paper, we propose a novel type of extremum-seeking controller, which continuously uses 1st-order least-squares fits to estimate the gradient of the performance map. The proposed method uses no dither signal, but utilizes a time window of history data of the performance map to estimate its gradient. The method is intuitive by nature, since the least-squares fit directly estimates the gradient, which can be visualized graphically. The approach allows for an asymptotic stability result, i.e. asymptotic convergence to the performance-optimal setting. Moreover, since the time-scale of the dither is eliminated, a possibly faster convergence speed of the algorithm can be obtained. Note that certain extremum-seeking schemes that reduce the dither amplitude in time, may yield asymptotic results as in e.g. [8], but at the expense of one more additional time-scale. The proposed extremum-seeking scheme utilizes only two tuning parameters (opposed to three or more in classical extremum-seeking schemes), which makes the scheme intuitive and easy to apply. We note that other continuous-adaptation based approaches have recently been suggested in literature which utilize observer-based schemes to estimate gradient (and Hessian) properties of the performance map, see e.g. [11], [3], [7], opposed to 1st-order least-squares fits using history data.

The main contributions of the paper can be summarized as follows. Firstly, we propose a novel two-parameter extremum-seeking scheme which is intuitive by nature, and easy to apply. Secondly, we present an asymptotic stability proof for the extremum-seeking scheme in combination with a static map. Thirdly, simulation results illustrate the effectiveness of the proposed extremum-seeking scheme, also for the dynamic setting.

The remainder of the paper is organized as follows. In Section II, we present the extremum-seeking approach, followed by an asymptotic stability proof for static maps in Section III. Simulation results are presented in Section IV. Conclusions and recommendations will be presented in Section V.

II. EXTREMUM-SEEKING CONTROL APPROACH USING LEAST-SQUARES FITS

Consider the extremum-seeking control scheme depicted in Fig. 1. The scheme consists of a stabilized plant, a performance function, a gradient estimator and an optimizer which adapts the parameter $\theta \in \mathbb{R}$. Additionally, two buffers

*This research is financially supported by the Dutch Technology Foundation STW.

¹B.G.B. Hunnekens, N. van de Wouw, and H. Nijmeijer are with the Eindhoven University of Technology, Department of Mechanical Engineering, P.O. Box 513, NL 5600 MB Eindhoven, The Netherlands B.G.B.Hunnekens@tue.nl

²M.A.M. Haring is with the Norwegian University of Science and Technology, Department of Engineering Cybernetics, 7491 Trondheim, Norway.

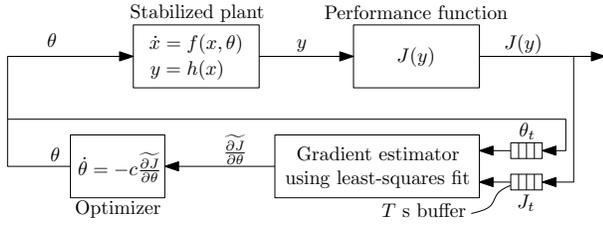


Fig. 1. Dither-free extremum-seeking scheme with least-squares fit of last T seconds to estimate gradient.

are present containing T seconds of history of the parameter $\theta_t \in \mathcal{C}$ and T seconds of history of the performance $J_t \in \mathcal{C}$, where $\mathcal{C} = \mathcal{C}([-T, 0], \mathbb{R})$ is the Banach space of continuous functions mapping the interval $[-T, 0]$ into \mathbb{R} . Moreover, $\theta_t \in \mathcal{C}$ is defined as $\theta_t(s) = \theta(t + s)$, for $-T \leq s \leq 0$ and J_t is defined as $J_t(s) = J(y(t + s))$.

The stabilized (possibly nonlinear) plant can be described by the following differential equation:

$$\begin{aligned} \dot{x} &= f(x, \theta) \\ y &= h(x), \end{aligned} \quad (1a) \quad (1b)$$

with state $x \in \mathbb{R}^n$, performance output $y \in \mathbb{R}$, and $\theta \in \mathbb{R}$ the performance parameter to be optimally tuned using the extremum-seeking scheme. For each fixed parameter θ , the plant is assumed to have a unique, globally asymptotically stable equilibrium point $x^*(\theta)$ (the assumptions will be made precise in Section III for the static-map setting).

The performance of the stabilized plant is characterized by the performance function $J(y) : \mathbb{R} \rightarrow \mathbb{R}$, see Fig. 1. The goal of the extremum-seeking controller is to minimize the *steady-state* performance map

$$J_{sta}(\theta) := J(y^*(\theta)) = J(h(x^*(\theta))), \quad (2)$$

which we assume to possess a unique minimum at θ^* , see Fig. 2.

The gradient estimator, see Figs. 1 and 2, continuously computes an estimate $\frac{\partial J}{\partial \theta}(\theta_t, J_t) : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}$ of the true gradient $\frac{\partial J_{sta}}{\partial \theta}(\theta(t))$, using a 1st-order least-squares fit of the last T seconds of data. The length T of the history interval used in the fit can be chosen by the designer. Details about the gradient estimation procedure are given in Section II-A. The estimated gradient $\frac{\partial J}{\partial \theta}(\theta_t, J_t)$ is used by the optimizer

$$\dot{\theta} = -c \frac{\partial J}{\partial \theta}(\theta_t, J_t), \quad (3)$$

in order to steer the parameter θ to its performance optimizing value θ^* , with adaptation gain $c > 0$, which can be chosen by the designer. Note that the closed-loop dynamics (1),(3) (with $J = J(y)$) is described by a functional differential equation (instead of an ordinary differential equation) due to the fact that the estimated gradient $\frac{\partial J}{\partial \theta}(\theta_t, J_t)$ depends on the last T seconds of history. Furthermore, note that the adaptation gain c and T are the only two parameters to be chosen by the designer of the extremum-seeking scheme, which simplifies the tuning procedure compared to classical

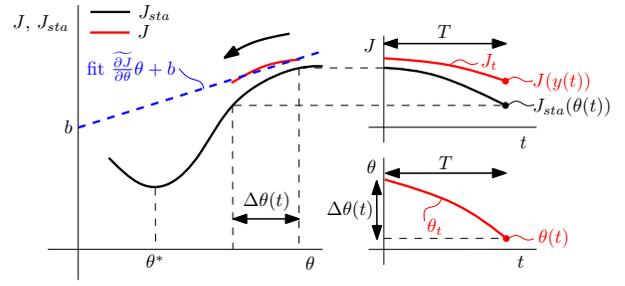


Fig. 2. Illustration of fit procedure in extremum-seeking controller.

extremum-seeking schemes, where typically more parameters should be chosen by the designer (such as low/high-pass filter parameters, and dither frequency and amplitude).

A. 1st-order least-squares gradient estimate

In contrast to classical extremum-seeking schemes, the gradient estimate in Fig. 1 is obtained without the use of an external dither signal applied to the parameter θ , see e.g. [6], [13]. In the extremum-seeking scheme proposed here, the last T seconds of data is used in order to continuously estimate the gradient using a 1st-order least-squares fit, see Fig. 2. The 1st-order least-squares fit $a\theta + b$ at time t follows from the following convex minimization problem

$$\min_{a,b} \int_{-T}^0 (J(y(t+\tau)) - (a\theta(t+\tau) + b))^2 d\tau. \quad (4)$$

Note that the explicit solution for the estimated gradient a can be computed in closed-form (or, in a digital implementation, a can also efficiently be computed by solving a set of normal equations [5, Section 4.6]).

As long as the adaptation gain c is chosen small enough to guarantee that the time-scale of the optimization is separated from (slower than) the time-scale of the plant-dynamics, the performance $J(y)$ will be close to the steady-state performance $J_{sta}(\theta) = J(y^*(\theta))$, i.e. the red line is close to the black line in Fig. 2, allowing for an accurate estimate a of the true gradient $\frac{\partial J_{sta}}{\partial \theta}(\theta(t))$. This is a classical time-scale separation result which is well-known in existing extremum-seeking literature, see e.g. [12]. Note that due to the absence of an external dither signal, there is however no need for time-scale separation between, on the one hand, the time-scale of the dither and, on the other hand, the optimizer and plant dynamics. This eliminates one of the time-scales present in classical dither-based extremum-seeking schemes, resulting in possibly faster convergence of the extremum-seeking controller.

The estimated gradient $\frac{\partial J}{\partial \theta}(\theta_t, J_t)$ used in the extremum-seeking loop is now given by:

$$\frac{\partial J}{\partial \theta}(\theta_t, J_t) = \begin{cases} a & \text{if } \Delta\theta(t) > 0 \\ 0 & \text{if } \Delta\theta(t) = 0, \end{cases} \quad (5)$$

with $\Delta\theta(t)$, see Fig. 2, defined as

$$\Delta\theta(t) := \underbrace{\max_{\tau \in [-T, 0]} \theta(t+\tau)}_{\theta_{max}(t)} - \underbrace{\min_{\tau \in [-T, 0]} \theta(t+\tau)}_{\theta_{min}(t)}. \quad (6)$$

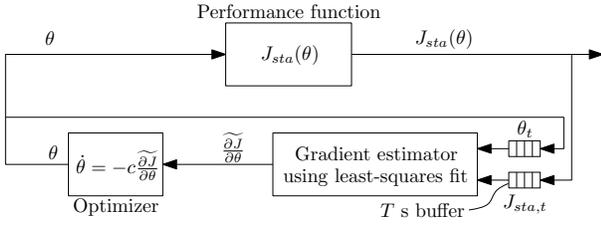


Fig. 3. Extremum-seeking scheme for static-map setting.

Note that if $\Delta\theta(t) = 0$, θ in fact resides in a single point in the time interval $[t - T, t]$, which cannot be used to make a 1st-order fit, hence we use $\frac{\partial \widetilde{J}}{\partial \theta}(\theta_t, J_t) = 0$ in (5) if $\Delta\theta(t) = 0$. For this reason, the extremum-seeking scheme needs to be initiated by a non-constant initial condition $\theta_0 \in \mathcal{C}$. This will be formalized in Section III, where we present a stability proof for the static-map setting.

III. STABILITY ANALYSIS FOR THE STATIC MAP SETTING

In this section, we will present the stability proof for the extremum-seeking scheme, in the static-map setting, see Fig. 3. Note that in the static-map setting there is no plant dynamics, such that in essence we are minimizing a static map $J_{sta}(\theta)$.

The dynamics of the extremum-seeking scheme for static maps is now governed by the following functional differential equation:

$$\dot{\theta}(t) = -c \frac{\partial \widetilde{J}}{\partial \theta}(\theta_t, J_{sta,t}), \quad (7)$$

where $J_{sta,t} \in \mathcal{C}$ is defined as $J_{sta,t} := J_{sta}(t + s)$ for $-T \leq s \leq 0$, and where $\frac{\partial \widetilde{J}}{\partial \theta}(\theta_t, J_{sta,t}) : \mathcal{C} \rightarrow \mathbb{R}$ is now defined as

$$\frac{\partial \widetilde{J}}{\partial \theta}(\theta_t, J_{sta,t}) = \begin{cases} a & \text{if } \Delta\theta(t) > 0, \\ 0 & \text{if } \Delta\theta(t) = 0, \end{cases} \quad (8)$$

with a given by the following ‘‘least-squares fit’’:

$$\min_{a,b} \int_{-T}^0 (J_{sta}(\theta(t+\tau)) - (a\theta(t+\tau) + b))^2 d\tau. \quad (9)$$

The following assumptions are made on the static performance map $J_{sta}(\theta)$.

Assumption III.1 The performance map J_{sta} and its first two derivatives with respect to θ are continuous and bounded on compact sets in θ . The map $J_{sta}(\theta)$ attains a unique minimum at $\theta = \theta^* \in \mathbb{R}$, and the following holds:

- $\frac{\partial J_{sta}}{\partial \theta}(\theta^*) = 0$;
- $\frac{\partial^2 J_{sta}}{\partial \theta^2}(\theta^*) > 0$;
- $\frac{\partial J_{sta}}{\partial \theta}(\theta)(\theta - \theta^*) > 0$ for all $\theta \in \mathbb{R} \setminus \theta^*$.

The following lemma provides a bound on the error of the gradient estimate, which is used later in the stability proof of Theorem 1.

Lemma III.2 Consider the extremum-seeking scheme for static maps as in (6)-(9), the assumptions on the static

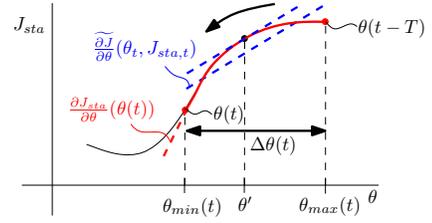


Fig. 4. Illustration of fit and error-bound in static-map setting.

map $J_{sta}(\theta)$ as in Assumption III.1, and initial conditions $\theta_0 \in \Theta_0 \subset \mathcal{C}$, where

$$\Theta_0 = \{\theta_0 \in \mathcal{C} \mid \max_{\tau \in [-T, 0]} |\theta(\tau) - \theta^*| \leq \rho_\theta, \Delta\theta(0) > 0\}, \quad (10)$$

with $\rho_\theta > 0$. Then, the error in the gradient estimate

$$\epsilon(\theta_t) = \frac{\partial \widetilde{J}}{\partial \theta}(\theta_t, J_{sta,t}) - \frac{\partial J_{sta}}{\partial \theta}(\theta(t)) \quad (11)$$

can be upper-bounded as follows:

$$|\epsilon(\theta_t)| \leq cT \max_{\tau \in [-T, 0]} \left| \frac{\partial^2 J_{sta}}{\partial \theta^2}(\theta(t+\tau)) \right| \max_{\tau \in [-2T, 0]} \left| \frac{\partial J_{sta}}{\partial \theta}(\theta(t+\tau)) \right|. \quad (12)$$

Proof: We only give a sketch of the proof here, a detailed proof is omitted for reasons of brevity.

Consider the situation sketched in Fig. 4. The true gradient at time t is given by $\frac{\partial J_{sta}}{\partial \theta}(\theta(t))$. The 1st-order least-squares fit results in an estimate $\frac{\partial \widetilde{J}}{\partial \theta}(\theta_t, J_{sta,t})$ as shown for example in Fig. 4. By the mean-value theorem it can be proven, and intuitively understood, that the estimated gradient $\frac{\partial \widetilde{J}}{\partial \theta}(\theta_t, J_{sta,t})$ is in fact equal to the true gradient $\frac{\partial J_{sta}}{\partial \theta}(\theta')$ for some $\theta' \in [\theta_{min}(t), \theta_{max}(t)]$, with $\theta_{min}(t) := \min_{\tau \in [-T, 0]} \theta(t+\tau)$ and $\theta_{max}(t) := \max_{\tau \in [-T, 0]} \theta(t+\tau)$. In other words, the estimated gradient at time t is equal to the true gradient at a time-instance during the last T seconds. Using this fact, see Fig. 4, it can be shown that the error $|\epsilon(\theta_t)|$ in the gradient estimate, can be upper-bounded by

$$|\epsilon(\theta_t)| \leq \max_{\tau \in [-T, 0]} \left| \frac{\partial^2 J_{sta}}{\partial \theta^2}(\theta(t+\tau)) \right| \Delta\theta(t). \quad (13)$$

Since $\theta(t)$ satisfies the dynamics in (7) and since the estimated gradient is equal to the true gradient somewhere during the last T seconds, it can be shown for $\Delta\theta(t)$ that

$$\Delta\theta(t) \leq cT \max_{\tau \in [-2T, 0]} \left| \frac{\partial J_{sta}}{\partial \theta}(\theta(t+\tau)) \right|, \quad (14)$$

such that, combined with (13), a bound on $\epsilon(\theta_t)$ as in (12) results. ■

The following theorem presents the main stability result of this paper.

Theorem 1 Consider the extremum-seeking scheme for static maps as in (6)-(9), with the assumptions on the static map $J_{sta}(\theta)$ as in Assumption III.1. Then, for any initial condition $\theta_0 \in \Theta_0 \subset \mathcal{C}$ with Θ_0 as in (10), and any $\rho_\theta > 0$,

there exists a product cT of the adaptation gain c and fit history time lapse T small enough such that

$$\lim_{t \rightarrow \infty} \theta(t) = \theta^* \quad \text{and} \quad \sup_{t \geq 0} |\theta(t) - \theta^*| \leq \rho_\theta. \quad (15)$$

Proof: First, decompose the estimated gradient in the following way,

$$\frac{\partial \tilde{J}}{\partial \theta}(\theta_t, J_{sta,t}) = \left(\frac{\partial J_{sta}}{\partial \theta}(\theta(t)) + \epsilon(\theta_t) \right), \quad (16)$$

with $\frac{\partial J_{sta}}{\partial \theta}(\theta(t))$ the true gradient at $\theta(t)$ and $\epsilon(\theta_t)$ an error term denoting the deviation between the true gradient and its estimate. Using this decomposition, the optimizer dynamics (7) can be written in the following way:

$$\dot{\theta}(t) = -c \left(\frac{\partial J_{sta}}{\partial \theta}(\theta(t)) + \epsilon(\theta_t) \right). \quad (17)$$

Note that $\theta = \theta^*$ is an equilibrium point of (6)-(9), since for constant $\theta = \theta^*$, $\Delta\theta = 0$. Hence, by (8), the estimated gradient $\frac{\partial \tilde{J}}{\partial \theta}(\theta_t, J_{sta,t}) = 0$, and, by (7), $\dot{\theta} = 0$. To investigate the stability of the optimal $\theta = \theta^*$, we consider the following candidate Lyapunov-Razumikhin function:

$$V(\theta(t)) = \frac{1}{2}(\theta(t) - \theta^*)^2 =: \frac{1}{2}\tilde{\theta}^2(t), \quad (18)$$

with $\tilde{\theta} := \theta - \theta^*$, such that we can write the time derivative \dot{V} along solutions of the system (17) as

$$\dot{V} = -c\tilde{\theta}(t) \left(\frac{\partial J_{sta}}{\partial \theta}(\theta(t)) + \epsilon(\theta_t) \right). \quad (19)$$

From Lemma III.2, it follows that we can bound the error in the gradient estimate $\epsilon(\theta_t)$ as

$$\begin{aligned} |\epsilon(\theta_t)| &\leq cT \max_{\tau \in [-T, 0]} \left| \frac{\partial^2 J_{sta}}{\partial \theta^2}(\theta(t + \tau)) \right| \\ &\quad \times \max_{\tau \in [-2T, 0]} \left| \frac{\partial J_{sta}}{\partial \theta}(\theta(t + \tau)) \right| \\ &\leq cTH \max_{\tau \in [-2T, 0]} \left| \frac{\partial J_{sta}}{\partial \theta}(\theta(t + \tau)) \right|, \end{aligned} \quad (20)$$

where we used that on compact sets $\theta \in \Theta$,

$$\Theta := \{\theta \in \mathbb{R} \mid |\theta - \theta^*| \leq \rho_\theta\}, \quad (21)$$

we can upper-bound the Hessian $\frac{\partial^2 J_{sta}}{\partial \theta^2}$ by $H := \max_{s \in \Theta} \left| \frac{\partial^2 J_{sta}}{\partial \theta^2}(s) \right|$. Later, we will conclude that for sufficiently small cT , Θ (which is a sub-level set of $V(\theta)$ in (18)) is a positive invariant set.

Substituting the bound (20) in (19), we obtain

$$\dot{V} \leq -c\tilde{\theta}(t) \frac{\partial J_{sta}}{\partial \theta}(\theta(t)) + c^2TH |\tilde{\theta}(t)| \max_{\tau \in [-2T, 0]} \left| \frac{\partial J_{sta}}{\partial \theta}(\theta(t + \tau)) \right|. \quad (22)$$

Using Assumption III.1, $\frac{\partial^2 J_{sta}}{\partial \theta^2}(\theta^*) > 0$, and $\frac{\partial J_{sta}}{\partial \theta}(\theta)(\theta - \theta^*) > 0 \forall \theta \in \Theta \setminus \theta^*$ imply that

$$\alpha_1 \tilde{\theta}^2 \leq \frac{\partial J_{sta}}{\partial \theta}(\theta) \tilde{\theta} \leq \alpha_2 \tilde{\theta}^2 \quad (23a)$$

$$\alpha_1 |\tilde{\theta}| \leq \left| \frac{\partial J_{sta}}{\partial \theta}(\theta) \right| |\tilde{\theta}|, \quad (23b)$$

for some $\alpha_2 \geq \alpha_1 > 0$. Using these bounds in (23) for \dot{V} in (22), we obtain:

$$\begin{aligned} \dot{V} &\leq -c\alpha_1 \tilde{\theta}^2(t) + c^2TH |\tilde{\theta}(t)| \max_{\tau \in [-2T, 0]} \left| \frac{\partial J_{sta}}{\partial \theta}(\theta(t + \tau)) \right| \\ &\leq -c\alpha_1 |\tilde{\theta}(t)|^2 + c^2TH |\tilde{\theta}(t)| \max_{\tau \in [-2T, 0]} \left(\alpha_2 |\tilde{\theta}(t + \tau)| \right) \end{aligned} \quad (24)$$

From the Lyapunov-Razumikhin theorem [2], a sufficient condition for stability is the existence of a continuous nondecreasing function $p(s) > s$ for $s > 0$ such that the following holds: $\dot{V}(t) < 0$ if the current value $p(V(\theta(t)))$ is larger than or equal to the one occurring in the last $2T$ seconds $\max_{\tau \in [-2T, 0]} V(\theta(t + \tau))$. Let $p(s) = d^2s$, with $d > 1$, then we require that $\dot{V}(t) < 0$ if $p(V(\theta(t))) \geq \max_{\tau \in [-2T, 0]} V(\theta(t + \tau))$. Hence, we require that $\dot{V}(t) < 0$ if

$$\frac{1}{2}d^2\tilde{\theta}^2(t) \geq \frac{1}{2} \max_{\tau \in [-2T, 0]} \left(\tilde{\theta}^2(t + \tau) \right), \quad (25)$$

or, equivalently, if

$$\max_{\tau \in [-2T, 0]} \left(\alpha_2 |\tilde{\theta}(t + \tau)| \right) \leq d\alpha_2 |\tilde{\theta}(t)|. \quad (26)$$

Using this in (24) gives:

$$\dot{V} \leq -c\alpha_1 |\tilde{\theta}(t)|^2 + c^2TH d\alpha_2 |\tilde{\theta}(t)|^2 \quad (27)$$

$$= -c\alpha_1 |\tilde{\theta}(t)|^2 \left(1 - cTH d \frac{\alpha_2}{\alpha_1} \right). \quad (28)$$

Note that $\dot{V} < 0$ for all $\tilde{\theta} \in \Theta \setminus \theta^*$ if $1 - cTH d \frac{\alpha_2}{\alpha_1} > 0$, or, equivalently, if

$$cT < \frac{\alpha_1}{Hd\alpha_2}, \quad (29)$$

such that $\lim_{t \rightarrow \infty} \theta(t) = \theta^*$ for all $\theta_0 \in \Theta_0$. Note that the set Θ , see (21) (which is a sub-level set of V) is a positively invariant set, hence the bound in (20) is valid, and hence $\sup_{t \geq 0} |\theta(t) - \theta^*| \leq \rho_\theta$. Choosing an 'initial condition' $\theta_0 \in \Theta_0$ as specified in (10), the Lyapunov-Razumikhin theorem guarantees that the Lyapunov function V can never become larger than the maximum V that occurred in the past. Note that we can always choose cT small enough in order to achieve the condition in (29). Moreover, note that for any arbitrarily large (though fixed) ρ_θ , and $d > 1$ the bounded values for α_1 , α_2 and H are fixed and we can choose the adaptation gain c and fit time T small enough in order to ensure that (29) is satisfied and hence $\lim_{t \rightarrow \infty} \theta(t) = \theta^*$. ■

The result in Theorem 1 is a type of asymptotic stability result, for initial conditions $\theta_0 \in \Theta_0$ that are non-constant (such that an initial fit can be made). In (15), the 'lim'-part relates to the attractive part of the asymptotic stability result and the 'sup'-part relates to the stability part of the asymptotic stability result.

Also note that it follows from Theorem 1 that only the product cT is important, which should be chosen small enough. Intuitively, a small product of c and T guarantees that a small neighborhood of $\theta(t)$ is used to obtain an accurate gradient estimate. In the static case, T can be chosen very small, allowing c to be chosen large (since

only product cT should be small), resulting in arbitrarily fast convergence. In the general case of a dynamical plant, as presented in Section II, we conjecture that, in addition, the adaptation gain c should be chosen small enough to guarantee that the performance $J(y)$ remains close to the steady-state performance $J_{sta}(\theta)$, see Fig. 2.

Remark III.3 In the extremum-seeking setting with dynamics, see Figs. 1 and 2, there will be an additional error term in the gradient estimate, due to the fact that the x -dynamics is not exactly in steady-state (i.e. $J(y) \neq J_{sta}(\theta)$). Therefore, it is no longer necessarily true that the estimated gradient is equal to the true gradient for some time instance during the last T seconds. However, we conjecture that for sufficiently small adaptation gain c (such that $J(y) \approx J_{sta}(\theta)$) and well-chosen fit-time T , an asymptotic stability result as in Theorem 1 is still possible, which is substantiated by the simulation results in Section IV. This is subject of future research.

IV. SIMULATION RESULTS

In this section, we will present simulation results in order to illustrate the effectiveness of the proposed extremum-seeking strategy. We will present results for the static-map setting as presented in Section III, and results on the dynamic setting as presented in Section II.

Consider the following simple dynamical plant of the form (1):

$$\dot{x} = -x + m(\theta) \quad (30a)$$

$$y = x, \quad (30b)$$

with $m(\theta)$ defined as $m(\theta) := 3 - \frac{1}{\sqrt{1+(\theta-2)^2}}$. Let the performance function $J = y$, from which it follows that the steady-state performance map is given by:

$$J_{sta}(\theta) = J(y^*(\theta)) = y^*(\theta) = x^*(\theta) = m(\theta), \quad (31)$$

where we used that $x^*(\theta) = m(\theta)$ (satisfying $\dot{x} = 0$). The map $J_{sta} = m(\theta)$ has a unique minimum at $\theta^* = 2$, see Fig. 5. Combined with the extremum-seeking scheme discussed in Section II, this describes the dynamic setting. The static-map setting simplifies to the scheme discussed in Section III, with $J_{sta}(\theta)$ as defined in (31).

A. The static-map setting

Since the least-squares gradient estimate requires T seconds of data before a fit can be made, we prescribe the evolution of θ for $t < 0$ for at least T seconds. Therefore, we use a prescribed $\dot{\theta} = P$ and do not use the adaptation in (7) for $t_0 \leq t \leq 0$. Note that this guarantees the initial condition $\theta_0 \in \Theta_0$ with Θ_0 as in (10). By choosing $t_0 \leq -T$, there is enough data available at $t = 0$ to estimate the gradient, and we switch on the adaptation according to (7). The used parameter values are shown in Table I.

The simulation result is shown in Fig. 5. The static performance map J_{sta} satisfies Assumption III.1. Therefore, from Theorem 1 it follows that there exist parameters c and T , with cT small enough such that the extremum-seeking controller

TABLE I
EXTREMUM-SEEKING PARAMETERS USED.

	Static setting	Dynamic setting
c	2	1
T	1	2
P	0.05	0.05
t_0	-1	-10
$\theta(t_0)$	5	5
$x(t_0)$	-	3

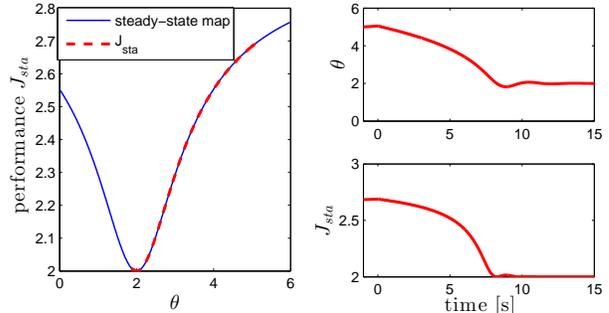


Fig. 5. Simulation result for extremum-seeking with fits, static-map setting.

converges asymptotically to the performance-optimal value $\theta^* = 2$. This is illustrated in Fig. 5. There is some overshoot at $\theta^* = 2$, which can be reduced by choosing the adaptation gain c smaller, or by choosing the fit-time T smaller. Note that choosing c smaller slows down the convergence, and that choosing T smaller results in a more accurate gradient estimate, since a ‘more local’ fit is made to estimate the gradient.

Remark IV.1 Related to the initial condition $\theta_0 \in \Theta_0$ in (10), also at the end of the simulation it is important that enough data is available to make a 1st-order fit (i.e. $\Delta\theta > 0$). Although theoretically it can be shown that the conditions in Theorem 1 guarantee that $\Delta\theta > 0$ for all $t \geq 0$, $\Delta\theta$ does become small when θ is converging towards θ^* . Therefore, it is recommended to only leave the adaptation on if the numerical conditioning of the least-squares fit is sufficiently good.

The estimated gradient $\frac{\partial \widetilde{J}}{\partial \theta}(\theta_t, J_{sta,t})$ and true gradient $\frac{\partial J_{sta}}{\partial \theta}(\theta(t))$ are shown in the upper-plot in Fig. 6. From this figure it is indeed apparent, as sketched in the proof of Lemma III.2, that the estimated gradient at time t is equal to the true gradient at some point in time during the last T seconds. Making the fit time T smaller will result in a better estimate of the estimated gradient.

B. The dynamic setting

The results for the dynamic setting, see Fig. 1, for plant dynamics (30), are shown in Fig. 7. The used parameter values can be found in Table I. Although no formal stability proof for the dynamic setting is presented in this paper, we see that the adaptation gain c and fit time T are chosen in such a way that θ converges asymptotically to the performance-optimal value θ^* . We also note that although the example shown here considers scalar x -dynamics, also plant-dynamics with $x \in \mathbb{R}^n$ can be considered, with $n \geq 2$.

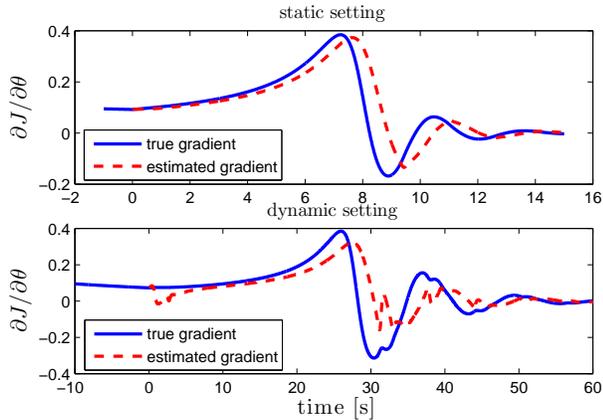


Fig. 6. Estimated gradients $\frac{\partial \tilde{J}}{\partial \theta}(\theta_t, J_{sta,t})$ (top) and $\frac{\partial \tilde{J}}{\partial \theta}(\theta_t, J_t)$ (bottom), and true gradients $\frac{\partial J_{sta}}{\partial \theta}(\theta(t))$ for static-map setting (top) and dynamic setting (bottom).

Note that since the x -dynamics is not in steady-state, the performance $J(y(t)) \neq J_{sta}(\theta(t))$. For $t \leq 0$, the adaptation according to (3) is turned off (since we prescribe $\dot{\theta} = P$), and the x -dynamics converges closely to its steady-state value $x^*(\theta) \approx 2.7$, see Fig. 7. Since c is chosen small enough, the performance $J(y)$ stays closely to the steady-state performance $J_{sta}(\theta)$ for $t > 0$, see Fig. 7.

A simulation result of a classical dither-based extremum-seeking scheme (see e.g. [6]), is also shown in Fig. 7. From this figure it is clear that θ converges to a neighborhood of θ^* , whereas the approach with fits, as presented in this paper, asymptotically converges to the performance optimal value θ^* . Since there is no dither signal, the performance does not oscillate around the steady-state performance, see Fig. 7. Because this time-scale of the dither is absent, this allows for a faster adaptation towards the optimum θ^* . Note that the neighborhood to which the dither-based approach converges can be made smaller, but at the expense of a slower convergence rate, which is a well-known tradeoff in extremum-seeking control, see e.g. [12].

The estimated gradient $\frac{\partial \tilde{J}}{\partial \theta}(\theta_t, J_t)$ and true gradient $\frac{\partial J_{sta}}{\partial \theta}(\theta(t))$ are shown in the bottom-plot in Fig. 6. Note that, opposed to the static setting, it is no longer necessarily true that the estimated gradient at time t is equal to the true gradient at some point in time during the last T seconds. For this reason, the authors believe that a (partially) different type of proof as the one presented for Theorem 1 should be considered in the dynamic case.

V. CONCLUSIONS

In this paper, we have proposed a new type of dither-free extremum-seeking controller which uses 1st-order least-squares fits to estimate the gradient of the performance map. History data of the performance map are used to estimate the gradient. The absence of dither allows for exact convergence of the performance-parameter to the optimum value. Moreover, the absence of dither eliminates one of the time-scales present in classical extremum-seeking schemes, allowing for a possibly faster convergence of the extremum-seeking algorithm. For the static-map setting a stability proof

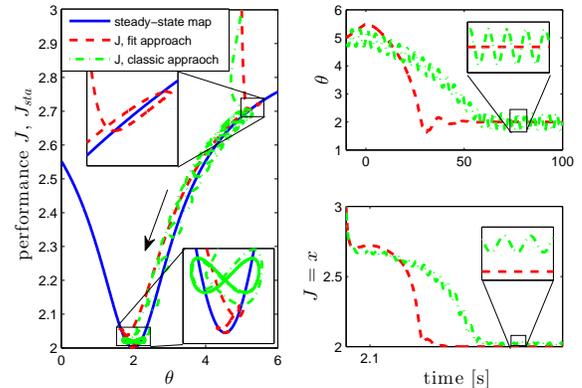


Fig. 7. Simulation results for extremum-seeking with fits (dashed) and of a classical dither-based extremum-seeking implementation (dash-dotted), dynamic setting.

has been presented which relies on a Lyapunov-Razumikhin type of proof for time-delay systems. Although the detailed stability proof is quite involved, the extremum-seeking approach is easy to apply and due to the fact that only two extremum-seeking parameters are used, also easy to tune. The presented method is also illustrated to work in case plant-dynamics is present. A formal stability proof of this dynamic setting is subject to future research.

REFERENCES

- [1] K.B. Ariyur and M. Krstić. *Real-time optimization by extremum-seeking control*. Wiley-Interscience, Hoboken, NJ, 2003.
- [2] K. Gu, V.L. Kharitonov, and J. Chen. *Stability of Time-Delay Systems*. Boston: Birkhäuser, 2002.
- [3] M. Guay, S. Dhaliwal, and D. Dochain. A time-varying extremum-seeking control approach. In *Proceedings of the American Control Conference*, pages 2643–2648, Washington DC, USA, 2013.
- [4] S.Z. Khong, D. Nešić, C. Manzie, and Y. Tan. Multidimensional global extremum seeking via the direct optimisation algorithm. *Automatica*, 49(7):1970–1978, 2013.
- [5] B. Kolman and D.R. Hill. *Elementary Linear Algebra*. Prentice Hall, Upper Saddle River, New Jersey, 2004.
- [6] M. Krstić and H.-H. Wang. Stability of extremum seeking feedback for general nonlinear dynamic systems. *Automatica*, 36(4):595–601, 2000.
- [7] W.H. Moase, C. Manzie, and M.J. Brear. Newton-like extremum-seeking part I: Theory. In *Proceedings of the IEEE Conference on Decision and Control*, pages 3839–3844, Shanghai, P.R. China, 2009.
- [8] S.J. Moura and Y.A. Chang. Lyapunov-based switched extremum seeking for photovoltaic power maximization. *Control Engineering Practice*, 21(7):971–980, 2013.
- [9] D. Nešić, T. Nguyen, Y. Tan, and C. Manzie. On a shubert algorithm-based global extremum seeking scheme. In *Proceedings of the IEEE Conference on Decision and Control*, pages 1615–1620, Maui, Hawaii, USA, 2012.
- [10] D. Nešić, Y. Tan, W.H. Moase, and C. Manzie. A unifying approach to extremum seeking: Adaptive schemes based on estimation of derivatives. In *2010 49th IEEE Conference on Decision and Control, CDC 2010*, pages 4625–4630, Atlanta, GA, USA, 2010.
- [11] J.J. Ryan and J.L. Speyer. Peak-seeking control using gradient and Hessian estimates. In *Proceedings of the 2010 American Control Conference, ACC 2010*, pages 611–616, Baltimore, MD, USA, 2010.
- [12] Y. Tan, W.H. Moase, C. Manzie, D. Nešić, and I.M.Y. Mareels. Extremum seeking from 1922 to 2010. In *Proceedings of the 29th Chinese Control Conference*, Beijing, China, 2010.
- [13] Y. Tan, D. Nešić, and I. Mareels. On non-local stability properties of extremum seeking control. *Automatica*, 42(6):889–903, 2006.
- [14] A.R. Teel and D. Popovic. Solving smooth and nonsmooth multi-variable extremum seeking problems by the methods of nonlinear programming. In *2001 American Control Conference*, volume 3, pages 2394–2399, Arlington, VA, 2001.