Automatica 95 (2018) 23–32

Contents lists available at ScienceDirect

Automatica

journal homepage: www.elsevier.com/locate/automatica

Brief paper

On the accuracy of gradient estimation in extremum-seeking control using small perturbations☆

Mark Haring *, Tor Arne Johansen

Centre for Autonomous Marine Operations and Systems, Department of Engineering Cybernetics, NTNU, the Norwegian University of Science and Technology, 7491 Trondheim, Norway

ARTICLE INFO

Article history:
Received 29 January 2015
Received in revised form 15 February 2018
Accepted 17 April 2018

Keywords:
Extremum-seeking control
Gradient estimation
Performance optimization
Stability
Nonlinear systems

ABSTRACT

In many extremum-seeking control methods, perturbations are added to the parameter signals to estimate derivatives of the objective function (that is, the steady-state parameter-to-performance map) in order to optimize the steady-state performance of the plant using derivative-based algorithms. However, large perturbations are often undesirable or inapplicable due to practical constraints and a high cost of operation. Yet, many extremum-seeking control algorithms rely solely on perturbations to estimate all required derivatives. The corresponding derivative estimates, especially the Hessian and higher-order derivatives, may be qualitatively poor if the perturbations are small. In this work, we investigate the use of the nominal parameter signals in addition to the perturbations to improve the accuracy of the gradient estimate. In turn, a more accurate gradient estimate may result in a faster convergence and may allow for a higher tuning-gain selection. In addition, we show that, if accurate curvature information of the objective function is available via estimation or a priori knowledge, it may be used to further enhance the accuracy of the gradient estimate.

1. Introduction

Extremum-seeking control is a collection of adaptive-control methods that optimize the steady-state performance of a plant in real time (Ariyur & Krstić, 2003; Khong, Nešić, Tan, & Manzie, 2013; Nešić, Tan, Manzie, Mohammadi, & Moase, 2012; Tan, Moase, Manzie, Nešić, & Mareels, 2010). By defining performance (or performance cost) as the output of a cost function of tunable plant parameters and measurable performance indicators, often no explicit knowledge of the plant dynamics is required. The steady-state relation between the parameters and the performance is commonly assumed to be given by a static input-to-output map (Krstić & Wang, 2000; Tan, Nešić, & Mareels, 2006), where the extremum of the map corresponds to the optimal steady-state performance. We refer to this map as the objective function. Many extremum-seeking control methods rely on extracting derivative information of the objective function from the parameters and performance signals of the plant (Nešić et al., 2012; Nešić, Tan, Moase, & Manzie, 2010). Subsequently, these derivatives are utilized by gradient-based (Krstić & Wang, 2000; Nešić et al., 2012; Tan et al., 2006) or Newton-based (Ghaffari, Krstić, & Nešić, 2012; Moase, Manzie, & Brear, 2010; Nešić et al., 2012) algorithms to steer the plant parameters towards the extremum of the objective function, thereby optimizing the steady-state plant performance. The majority of extremum-seeking control methods utilize perturbations to ensure that the parameter signals are sufficiently rich to estimate the derivatives of the objective function. The derivative estimates are obtained by correlating the perturbations and the time signal of the plant performance (Atta, Johansson, & Gustafsson, 2015; Eveleigh, 1967; Krstić & Wang, 2000; Poveda & Teel, 2017; Tan et al., 2006). The nominal part of the parameter signals is often ignored. The true values of the derivatives are commonly not obtained, because the performance of the plant is unequal to the steady-state performance due to plant dynamics and measurement noise.

To keep the dynamic transients of the performance signal small, the controller is generally chosen to be slower than the dynamics of the plant (Krstić & Wang, 2000; Moase, Tan, Nešić, & Manzie, 2011). For a limited class of plants, high-amplitude high-frequency perturbations can be used to overpower the original plant dynamics, enforcing an arbitrarily fast convergence upon the plant (Dürr, Stanković, Ebenbauer, & Johansson, 2013; Moase & Manzie, 2012; Scheinker & Krstić, 2017; Zhang, Siranosian, & Krstić, 2007).

☆ This work is a partly sponsored by the Research Council of Norway by the KMB project D2V, project number 210670, and through the Centres of Excellence funding scheme, project number 223254 (NTNU-AMOS). The material in this paper was not presented at any conference. This paper was recommended for publication in revised form by Associate Editor Raul Ordonez under the direction of Editor Miroslav Krstic.

* Corresponding author.
E-mail addresses: mark.haring@ntnu.no (M. Haring), tor.arne.johansen@ntnu.no (T.A. Johansen).

https://doi.org/10.1016/j.automatica.2018.05.001
0005-1098/© 2018 Elsevier Ltd. All rights reserved.

However, contrary to these highly invasive methods, in practice, one often wishes to keep the disruption of nominal operation to a minimum to keep operational costs low, the response of the plant predictable, and state and output values within predefined limits. This can be achieved by using small perturbations. The use of small perturbations results in relatively little perturbation-related content in the time signal of the plant performance, which may lead to poor estimates of the derivatives of the objective function, especially in the presence of measurement noise. This is particularly true for the estimate of the Hessian and higher-order derivatives of the objective function. Therefore, gradient-based algorithms may be preferred over Newton-based ones in this case. Additionally, the contribution of the nominal part of the parameter signals to the performance signal of the plant is relatively large if the perturbations are small. With the help of an observer, the nominal part of the parameter signals may be included in the estimation process to increase the accuracy of the derivative estimates (Clark & Godfrey, 1966; Gelbert, Moeck, Paschereit, & King, 2012; Guay & Dochain, 2015; Ryan & Speyer, 2010). Although not specifically identified as such, the effect of the nominal part of the parameter signals on the produced derivative estimates is investigated in Gelbert et al. (2012) and Guay and Burns (2014) by a comparison of various extremum-seeking control methods. However, due to the significant differences between the used methods, it is unclear if the observed results are due to the use of the nominal part of the parameter signals or due to any other structural difference. Moreover, because extremum-seeking control is highly dependent on tuning, the obtained performance of any extremum-seeking method is for a large extent determined by the tuning capabilities of the user.

In this work, we introduce an extremum-seeking controller for which the contribution of the nominal part of the parameter signals to the gradient estimate can be isolated. Therefore, the effect of incorporating the nominal parameters in the estimation process can be investigated using a single controller, which largely eliminates the challenges that affect the comparisons in Gelbert et al. (2012) and Guay and Burns (2014). In addition, we show that curvature information of the objective function, if available, may further enhance the accuracy of the gradient estimate. The results in this work may be regarded as an extension of the results in Haring and Johansen (2017) in which the nominal parameters and curvature information are not utilized for gradient estimation.

This work is organized as follows. After introducing the extremum-seeking problem in Section 2, our controller is presented in Section 3. A stability analysis of the closed-loop optimization scheme is provided in Section 4. In Section 5, we study, with the help of simulation examples, the effects of incorporating in the gradient estimate the nominal parameter signals and curvature information of the objective function. The conclusion of this work is presented in Section 6. The sets of real numbers, positive real numbers and nonnegative real numbers are denoted by $\mathbb{R}$, $\mathbb{R}_{>0}$ and $\mathbb{R}_{\geq 0}$, respectively. The sets of natural numbers (nonnegative integers) and positive integers are given by $\mathbb{N}$ and $\mathbb{N}_{>0}$. The Euclidean norm is denoted by $\|\cdot\|$. We write the identity matrix and zero matrix as $\mathbf{I}$ and $\mathbf{0}$, respectively.

2. Formulation of the extremum-seeking problem

Consider the following multi-input-single-output nonlinear system:

$$\begin{aligned} \dot{\mathbf{x}}(t) &= \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t)) \\ y(t) &= h(\mathbf{x}(t), \mathbf{u}(t)) + d(t), \end{aligned} \tag{1}$$

where $\mathbf{x} \in \mathbb{R}^{n_\mathbf{x}}$ is the state, $\mathbf{u} \in \mathbb{R}^{n_\mathbf{u}}$ is the input, $y \in \mathbb{R}$ is the output, $d \in \mathbb{R}$ is a disturbance and $t \in \mathbb{R}_{\geq 0}$ is the time. The dimensions of the state and the input are given by $n_\mathbf{x}, n_\mathbf{u}, \in \mathbb{N}_{>0}$, respectively. In the context of extremum-seeking control, the system can be regarded as a cascade of the plant and the cost function that quantifies the performance of the plant (see Haring, van de Wouw, & Nešić, 2013 for example), where the input $\mathbf{u}$ is a vector of tunable plant parameters and the output y is the output of the cost function, which we call the performance measurement. The output of the function h is the output of the cost function in the absence of measurement noise. The disturbance d represents the contribution of measurement noise to the output of the cost function. The state $\mathbf{x}$, the disturbance d and the functions $\mathbf{f}$ and h are unknown. Therefore, the relation between the parameters and the performance of the plant is unknown.

We make several assumptions with respect to the input-to-output behavior of the system in order to optimize the steady-state performance of the plant. First, we assume that, for each constant vector of plant parameters $\mathbf{u}$, there exists a constant steady-state solution of the system denoted by $\mathbf{x} = \mathbf{X}(\mathbf{u})$.

Assumption 1. There exist a twice continuously differentiable map $\mathbf{X} : \mathbb{R}^{n_\mathbf{u}} \to \mathbb{R}^{n_\mathbf{u}}$ and a constant $L_\mathbf{X} \in \mathbb{R}_{>0}$ such that

$$\mathbf{0} = \mathbf{f}(\mathbf{X}(\mathbf{u}), \mathbf{u}), \qquad \left\| \frac{d\mathbf{X}}{d\mathbf{u}}(\mathbf{u}) \right\| \leq L_\mathbf{X} \tag{2}$$

for all $\mathbf{u} \in \mathscr{U}$.

In addition, we assume that the steady-state solution $\mathbf{x} = \mathbf{X}(\mathbf{u})$ is unique and exponentially stable for constant inputs.

Assumption 2. There exist constants $\mu_\mathbf{x}, \nu_\mathbf{x} \in \mathbb{R}_{>0}$ such that, for each constant $\mathbf{u} \in \mathbb{R}^{n_\mathbf{u}}$, the solutions of the system satisfy

$$\|\tilde{\mathbf{x}}(t)\| \leq \mu_\mathbf{x} \|\tilde{\mathbf{x}}(t_0)\| e^{-\nu_\mathbf{x}(t-t_0)}, \tag{3}$$

with $\tilde{\mathbf{x}}(t) = \mathbf{x}(t) - \mathbf{X}(\mathbf{u})$, for all $t \geq t_0 \geq 0$ and all $\mathbf{x}(t_0) \in \mathbb{R}^{n_\mathbf{x}}$.

The disturbance-free steady-state relation between constant plant parameters and the plant performance can now be expressed by the static input-to-output map

$$F(\mathbf{u}) = h(\mathbf{X}(\mathbf{u}), \mathbf{u}). \tag{4}$$

We refer to the map F as the objective function. We assume that the cost function is designed such that there exists a unique minimum of the objective function that corresponds to the optimal steady-state plant performance. This is formulated as follows.

Assumption 3. The objective function $F : \mathbb{R}^{n_\mathbf{u}} \to \mathbb{R}$ is twice continuously differentiable and contains a unique minimum. Let $\mathbf{u}^*$ denote the corresponding minimizer. There exist constants $L_{F1}, L_{F2} \in \mathbb{R}_{>0}$ such that

$$\frac{dF}{d\mathbf{u}}(\mathbf{u})(\mathbf{u} - \mathbf{u}^*) \geq L_{F_1} \|\mathbf{u} - \mathbf{u}^*\|^2, \qquad \left\| \frac{d^2F}{d\mathbf{u}d\mathbf{u}^T}(\mathbf{u}) \right\| \leq L_{F2} \tag{5}$$

for all $\mathbf{u} \in \mathbb{R}^{n_\mathbf{u}}$.

Although the exact formulation may vary, assumptions on the existence and the attractiveness of the steady-state solution of the system, and the existence of an extremum[1] of the objective function are common in extremum-seeking control (Guay & Dochain, 2015; Krstić & Wang, 2000; Tan et al., 2010, 2006). To guarantee the soundness of the stability analysis in Section 4, we assume in addition that the following bounds on the derivatives of the functions $\mathbf{f}$ and h hold.

[1] The extremum is a minimum in this case.