

An Introduction to Extremum-Seeking Control

FRTN15 *Predictive and Adaptive Control*

LTH, Sweden

Alexander Skafte tfy13ask@student.lu.se

April, 2017

Contents

1	Introduction	2
2	Theory	3
3	Example and simulation results	6
4	Conclusion	9

1 Introduction

While common methods for adaptive control, both linear and nonlinear, aim to drive the output of a process to a known set point or reference trajectory, the goal of *extremum-seeking control* is to find a system input such that some measure of the process output is held at an extremum point, i.e. a minimum or a maximum. Different methods of achieving this goal have been studied and developed over the last century, and today there exist a multitude of such methods; some are discussed in [4] and [5]. In this paper we will investigate a simple perturbation-based scheme for performance optimization of a nonlinear dynamic plant. Under certain assumptions such a system can be shown to be stable; a proof for this is given by Krstić and Wang in [1].

Consider the following system:

- An unknown nonlinear dynamical plant $\dot{x} = f(x, u)$
- An unknown measured performance function $y = g(x)$
- Plant state $x \in \mathbb{R}^n$
- Plant input $u \in \mathbb{R}$
- $f : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ and $g : \mathbb{R}^n \rightarrow \mathbb{R}$ are both smooth

If we know a smooth state-feedback control law $u = \alpha(x, \theta)$, $\theta \in \mathbb{R}$ then the equilibria of the closed-loop system $\dot{x} = f(x, \alpha(x, \theta))$ is parameterized by θ , i.e. there is a θ which optimizes y .

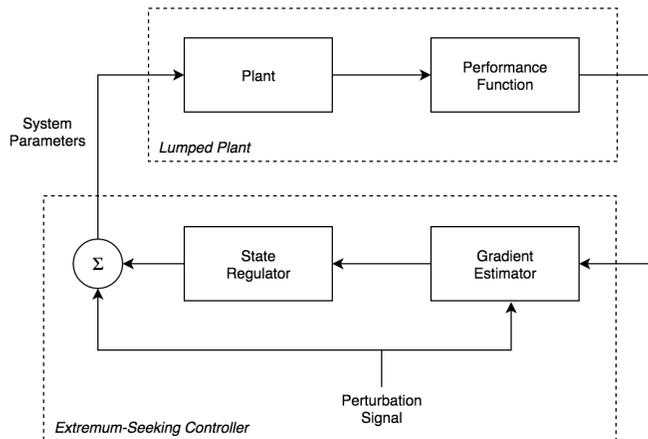


Figure 1: A general extremum-seeking scheme.

The basic idea behind the scheme investigated in this paper is to perturb the system with a slow periodic signal commonly chosen to be sinusoidal; by then

observing and acting on the output of the performance function one can estimate its gradient, which can then be used by the state regulator, commonly chosen to be an integrator, to iteratively seek an extremum point. In Section 2 we will first look at the case where we assume plant dynamics to be driven to stationarity by a control law and thus view the reference-to-output map as a static map. Thereafter we will briefly look at the case where the plant dynamics have not been driven to stationary, and in Section 3 observe by simulations how it affects performance with respect to a static performance function.

2 Theory

In this section we will develop an intuitive description of a standard perturbation-based extremum-seeking control scheme. First, assume the reference-to-output map to be static. The scheme passes the output of the performance function through a high-pass filter and thereafter multiplies the result by the perturbation signal in order to create a rough estimate of the gradient. In some cases this signal is thereafter also passed through a low-pass filter, which may improve the system slightly by for instance attenuating measurement noise. We will however not consider such a low-pass filter in the investigated scheme, since we aim to keep the analysis simple and merely provide an intuitive explanation of the principle. For a more rigorous treatment of extremum-seeking control, one may refer to [2, 3, 4] or any other literature on the subject.

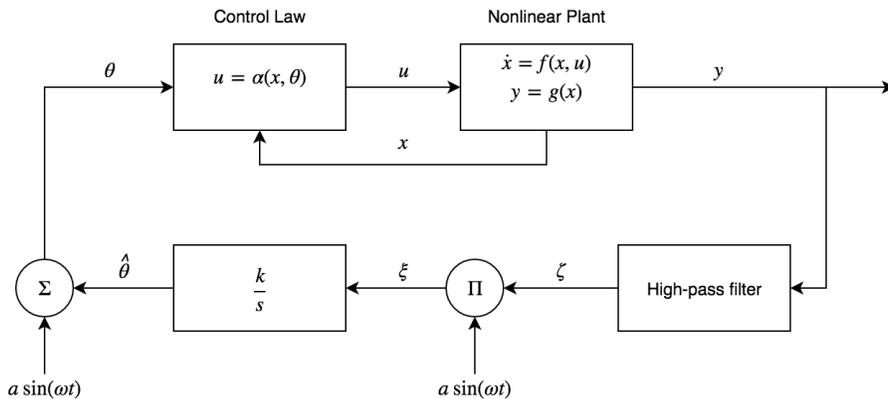


Figure 2: An extremum-seeking scheme.

In our extremum-seeking control system we have as input to the plant the sum of the estimated parameter $\hat{\theta}$ and a sinusoidal perturbation signal, i.e.

$$\theta = \hat{\theta} + a \sin(\omega t)$$

where a and ω are design parameters. The amplitude a provides a trade-off between speed of convergence and narrowness of region of attraction; a small a increases the probability of getting stuck at a local extremum but decreases the residual error at the extremum reached by the algorithm, while a large a increases speed of convergence but also the residual error. The perturbation frequency ω controls time scale separation of the parameter estimation process, conducted by the state regulator, and the gradient estimation process performed by the inclusion of a perturbation signal; a higher ω gives a cleaner gradient estimate, but the plant may be sensitive to high-frequency oscillations in which case a high ω is undesirable.

Since we assume the plant has reached steady-state we can view the plant as a static map from θ to its output y , i.e.

$$y = g(\theta) = g(\hat{\theta} + a \sin(\omega t))$$

We now wish to informally show that an extremum-seeking scheme using gradient estimation can be developed with only knowledge of θ and the output of the plant. Let ∂_x denote the (partial) differential operator with respect to x . Since θ is only slowly time-varying due to the frequency of the perturbation signal being chosen low we can approximate g at $\hat{\theta}$ using a first-degree Taylor expansion $f_a(x) = f(a) + \partial_x f(a)(x - a)$:

$$\begin{aligned} g_{\hat{\theta}}(\theta) &= g(\hat{\theta}) + \partial_{\theta} g(\hat{\theta})(\hat{\theta} + a \sin(\omega t) - \hat{\theta}) \\ &= g(\hat{\theta}) + \partial_{\theta} g(\hat{\theta})a \sin(\omega t) \end{aligned}$$

Here we have an expression where one term is an approximation of the gradient of g , apart from a slowly time-varying sinusoidal factor. This is a positive result; extracting the gradient by means of high-pass filtering the signal, which removes the DC-component $g(\hat{\theta})$, means that we are one step closer to being able to use the approximate gradient $\partial_{\theta} g(\hat{\theta})$ to seek an extremum point of g .

Consider a high-pass filter \mathcal{H} with a cutoff frequency ω_h chosen such that $\omega_h < \omega$ yet sufficiently large to remove the DC-component. Denoting the high-pass filtered approximation $g_{\hat{\theta}}(\theta)$ of $g(\theta)$ by ζ we have

$$\begin{aligned} \zeta &= \mathcal{H}\{g_{\hat{\theta}}(\theta)\} \\ &= \mathcal{H}\{g(\hat{\theta}) + \partial_{\theta} g(\hat{\theta})a \sin(\omega t)\} \\ &\approx \partial_{\theta} g(\hat{\theta})Aa \sin(\omega t + \varphi) \end{aligned}$$

where A and φ are due to the filter. Since $\omega_h < \omega$, A and φ have the following properties (proof omitted; however, the result can be seen by observing amplitude and phase diagrams of a general first-order high-pass filter, which we assume is used):

- (i) $\frac{1}{\sqrt{2}} < A < 1$
- (ii) $0 < \varphi < \frac{\pi}{4}$

Multiplying ζ by the perturbation signal $a \sin(\omega t)$ gives us the following expression:

$$\begin{aligned}\xi &= a \sin(\omega t) \zeta \\ &= a \sin(\omega t) \partial_{\theta} g(\hat{\theta}) A a \sin(\omega t + \varphi) \\ &= a^2 A \sin(\omega t) \sin(\omega t + \varphi) \partial_{\theta} g(\hat{\theta})\end{aligned}$$

Given that $\omega_h < \omega$ it can be shown that the product $\sin(\omega t) \sin(\omega t + \varphi)$ is positive "most of the time"; for $h \rightarrow \omega^+$ the product is positive approximately 94% of the time and for $\omega_h = \omega/2$ the product is positive approximately 99% of the time (this can be seen instantly by observing the product's graph).

In order to simplify the calculations and improve understanding of the method we simplify the above expression by assuming that the product of sinusoids is always positive and multiply this product with the constants $a^2 A$, giving us the expression

$$\begin{aligned}S(t) &= a^2 A \sin(\omega t) \sin(\omega t + \varphi) \\ \xi &= S(t) \cdot \partial_{\theta} g(\hat{\theta}), \quad 0 \lesssim S(t) \lesssim 1\end{aligned}$$

Hence ξ is equal to the estimated gradient of g at $\hat{\theta}$ except for the time-varying factor $S(t)$. This means that we can apply gradient-seeking in order to find a $\hat{\theta}$ which maps to an extremum point of g . We can update $\hat{\theta}$ by means of an integrator. Intuitively it is easy to see this if we consider a discrete-time zero-order hold integrator

$$I_{\tau,k}(z) = \frac{\tau k}{z - 1}$$

where τ is the sampling time and k , the integrator gain, is positive if we are seeking a maximum or negative if we are seeking a minimum; applying this integrator to ξ_i (the value of ξ at iteration i) we get an update of $\hat{\theta}$ of the form

$$\hat{\theta} = \frac{\tau k}{z - 1} \xi_i \quad \Leftrightarrow \quad \hat{\theta}_{i+1} = \hat{\theta}_i + \tau k \xi_i.$$

or a corresponding expression for a continuous-time integrator $I_k(s) = \frac{k}{s}$. Since $\tau > 0$ and $\xi \geq 0$ the sign of the $\tau k \xi$ term depends only on the sign of k , which we said to be positive if we are seeking a maximum and negative if we are seeking a minimum. As the term $\tau k \xi$ guides $\hat{\theta}$ toward an extremum point of g the gradient $\partial_{\theta} g(\hat{\theta})$ will eventually reach and oscillate around 0, which implies that an extremum of y has been found (the oscillation is due to the sinusoidal perturbation signal).

In summary we have, for an integrator \mathcal{I} ,

$$\begin{aligned} g(\theta) &\approx g(\hat{\theta}) + \partial_{\theta}g(\hat{\theta})a \sin(\omega t) \\ \zeta &\approx \partial_{\theta}g(\hat{\theta})Aa \sin(\omega t + \varphi) \\ \xi &= a \sin(\omega t)\zeta \gtrsim 0 \\ \hat{\theta} &= \mathcal{I}\{\xi\} \\ \theta &= \hat{\theta} + a \sin(\omega t) \end{aligned}$$

3 Example and simulation results

In this section we will first investigate how the parameters of the extremum-seeking controllers affects performance. For purposes of testing we will use the performance function seen in figure 3 which has local minimum at $g(-1) = 1$ and global minimum at $g(1) = -3$.

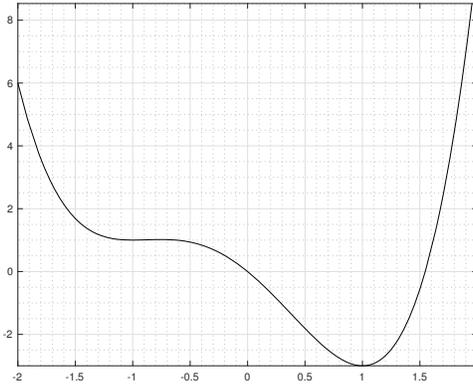


Figure 3: The performance function $g(\theta) = \theta^4 + \theta^3 - 2\theta^2 - 3\theta$.

First we look at the case where the reference-to-output map is static and thus described only by $g(\theta)$, the function in figure 3. As design parameters of the extremum-seeking controller we have perturbation amplitude a and frequency ω , integrator gain k and high-pass filter cutoff frequency ω_h . In figures 4–11 we see simulations for different parameter values of a , ω and k . For all simulations we use $\omega_h = 1$ as the high-pass filter cutoff frequency and $\theta = -1$ as an initial guess of θ (blue in the figures below).

Comparing figures 4–6 to figure 7 we see that a larger perturbation signal amplitude a makes the algorithm converge faster. However, as evident by inspection

of figures 8-11, we see that if a is too small the algorithm will not converge at all, due to its inability to "escape" the local minimum at $g(-1) = 1$.

A larger integrator gain k also decreases convergence time but causes a transient spike. The perturbation signal frequency ω only has a slight effect on convergence time, a result which is due to the quality of the gradient estimate obtained by varying ω .

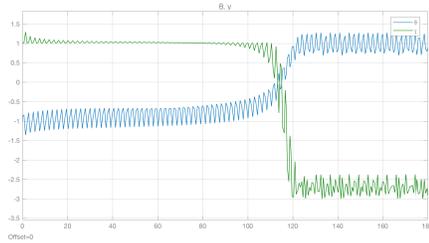


Figure 4: $a = 0.3, \omega = 3, k = -1$

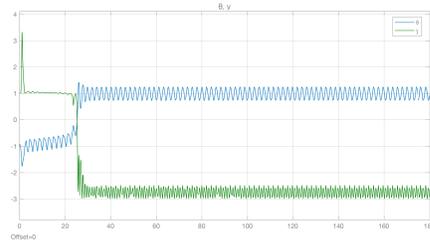


Figure 5: $a = 0.3, \omega = 3, k = -5$

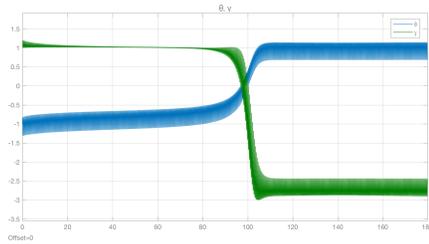


Figure 6: $a = 0.3, \omega = 15, k = -1$

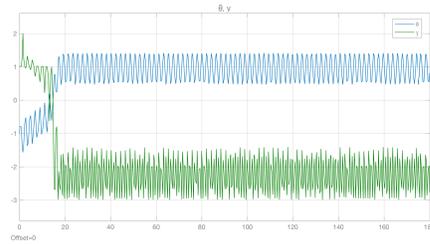


Figure 7: $a = 0.5, \omega = 3, k = -1$

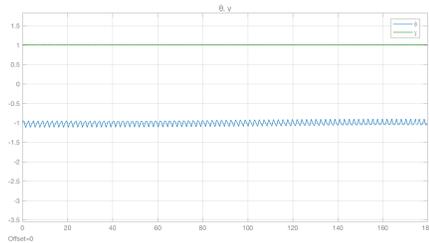


Figure 8: $a = 0.1, \omega = 3, k = -1$

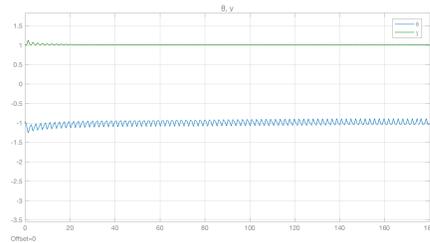


Figure 9: $a = 0.1, \omega = 3, k = -5$

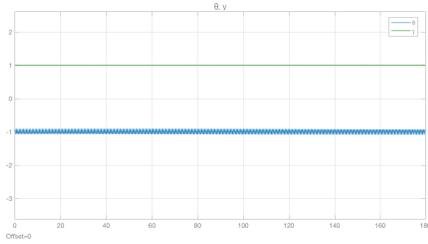


Figure 10: $a = 0.1, \omega = 10, k = -1$

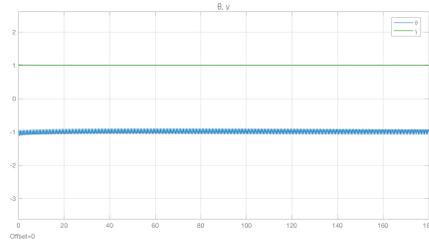


Figure 11: $a = 0.1, \omega = 10, k = -5$

Next we consider a system where the map from reference to output is not static; instead, θ is first passed through a dynamic system \mathcal{G} , the output of which is used as input to the performance function giving us $y = g(\mathcal{G}\{\theta\})$ as opposed to just $y = g(\theta)$. Relating the new controller to the one in figure 2 we just consider the control law to be $u = \theta$.

We can, as an example, let \mathcal{G} be described by the transfer function $G(s) = 1/(\tau s + 1)$ which makes it a low-pass filter with a cutoff frequency $\omega_l = 1/\tau$. By intuitive reasoning we conclude that too high a frequency of the perturbation signal will make finding the global extremum slow or even impossible.

The above reasoning is easily confirmed by simulations, as seen in figure 12 and 13. The parameters used for simulation were $\tau = 0.1, a = 0.3, \omega_h = 1$ and $k = -1$; figure 12 uses perturbation frequency $\omega = 3$ and figure 13 uses perturbation frequency $\omega = 15$. We see that for $\omega = 3$ the extremum-seeking controller manages to find the minimum of y but for $\omega = 15$ it does not due to the low pass filter removing too much of the perturbation signal.

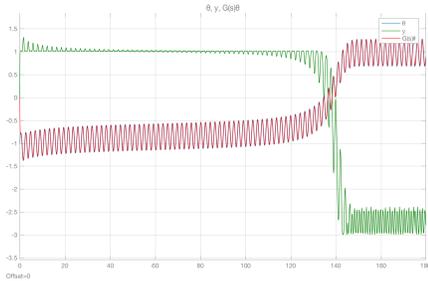


Figure 12: $\omega = 3$

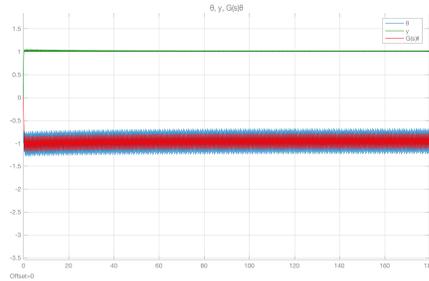


Figure 13: $\omega = 15$

4 Conclusion

Extremum-seeking control is a means of optimizing some measure of the output of a process with respect to an input parameter that drives the system. There exist a multitude of extremum-seeking control scheme, the perturbation-based one being the focus of this report. The inclusion of a perturbation signal allows one to estimate the gradient of the process output and, using an integrator, update the input parameter such that the process output is driven toward an extremum point. There are several design parameters such as perturbation signal amplitude and frequency, filter cutoff frequencies and integrator gain, each being a trade-off between different qualities within the system such as convergence time, stability, transient behaviour and so on. Should the process dynamics be difficult to efficiently stabilize with the control law, one has to take additional care in choosing the frequency of the perturbation signal such that the signal is not undesirable attenuated or amplified.

References

- [1] M. Krstic and H.H. Wang. *Stability of extremum seeking feedback for general dynamic systems*. *Automatica*, 36 (4):595–601, 2000.
- [2] M. Krstic and K.B. Ariyur. *Real-Time Optimization by Extremum-Seeking Control*. Hoboken, New Jersey: John Wiley & Sons, Inc., 2003.
- [3] M. Krstic and S.J. Liu. *Stochastic Averaging and Stochastic Extremum Seeking*. London: Springer, 2012.
- [4] C. Zhang and R. Ordóñez. *Extremum-Seeking Control and Applications: A Numerical Optimization-Based Approach*. London: Springer, 2012.
- [5] B. Calli, W. Caarls, P. Jonker and M. Wisse. *Comparison of Extremum Seeking Control Algorithms for Robotic Applications*. *Intelligent Robots and Systems*, IEEE, 2012.