Delft University of Technology

Document Version
Final published version

Licence
CC BY

Citation (APA)
Mulders, S. P., Rotea, M. A., & Gallo, A. J. (2026). Convergence of extremum seeking control for state-dependent and bilinear objectives. *Systems and Control Letters*, *217*, Article 106563. https://doi.org/10.1016/j.sysconle.2026.106563

Important note
To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright
In case the licence states “Dutch Copyright Act (Article 25fa)”, this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse
Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy
Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Systems & Control Letters 217 (2026) 106563

Contents lists available at ScienceDirect

Systems & Control Letters

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

Convergence of extremum seeking control for state-dependent and bilinear objectives

Sebastiaan Paul Mulders [a,*], Mario A. Rotea [b], Alexander Julian Gallo [c]

[a] Delft University of Technology, Netherlands
[b] The University of Texas at Dallas, United States of America
[c] University College London, United Kingdom

ARTICLE INFO

Keywords:
Extremum seeking control
Wiener-type systems
Objective function analysis
Convergence
Linear analysis

ABSTRACT

Extremum seeking control (ESC) is a model-free and adaptive scheme for optimizing steady-state performance of dynamic systems in real time, particularly when system models are unavailable or uncertain. Classical ESC relies on exciting the system with a sinusoidal dithering signal, thereby ensuring convergence to the optimum, provided that the dither excitation frequency remains well below the dominant system dynamics. However, this time-scale separation constraint limits convergence speed and robustness, especially in systems whose dominant dynamics begin at low frequencies. Enhanced ESC methods, such as proportional-integral schemes, phasor-based estimation techniques, and frequency-domain approaches, address the time-scale separation limitation. However, the understanding of how objective definitions with identical steady-state optima affect convergence remains limited. This paper analyzes ESC closed-loop convergence in dynamical systems for state-dependent and bilinear (state-input-dependent) objectives using a Wiener-type model framework. Although both objectives share equal steady-state optima, they exhibit fundamentally different convergence behavior when the dither frequency is chosen beyond the system's dominant dynamics (fast time scale). New insights are obtained through sequential linearization combined with frequency-domain examinations and signal-based analyses. Furthermore, improvements to the classical ESC scheme for dynamical bilinear objectives are proposed to relax time-scale separation constraints, enabling higher dither frequencies and leading to consistent and accelerated convergence.

1. Introduction

Extremum seeking control (ESC) finds its origin in the 1950s when the first papers on the algorithm emerged [1,2]. Initially, ESC was primarily used as a framework for optimizing the steady-state performance of slowly time-varying dynamical systems in uncertain environments. Over the decades, ESC has evolved as a powerful excitation-based optimization algorithm, capable of addressing nonlinear, time-varying, and multivariable systems, especially in applications where system models are either unavailable or highly uncertain [3,4]. After a period of reduced interest in the algorithm, ESC made a comeback in the early 2000s, when rigorous stability analyses and proofs were presented based on averaging and singular perturbation theory [5–7], formalizing and extending the early heuristic assessments. These theoretical advancements enabled ESC to be applied to complex real-world systems such as fuel cells [8], bioreactors [9], optical communication networks [10], and, more recently, for accelerating reinforcement learning policy optimization [11].

Classical extremum seeking control is a model-free and adaptive control algorithm that dynamically searches for the input that (locally) optimizes an online measurable performance objective [6]. It assumes that the system under optimization exhibits an optimum in its static (i.e., steady-state) input–output mapping. The ESC algorithm treats the steady-state (or slowly time-varying) part of the system input as a decision variable and adjusts it in real-time to find the optimum of the considered objective [12]. ESC typically employs periodic excitation, filtering, and signal demodulation techniques to estimate gradients and facilitate gradient-descent-like adaptation without requiring an explicit dynamic system model.

Classical single-input single-output (SISO) ESC schemes rely on the principles of time-scale separation, which can be categorized into three time scales: fast, medium, and slow. The fast scale corresponds to the inherent system dynamics, and the slow scale is reserved for system variations [13]. The medium time scale in classical ESC implementations is where the dither frequency should be chosen to follow

* Corresponding author.
E-mail address: s.p.mulders@tudelft.nl (S.P. Mulders).

https://doi.org/10.1016/j.sysconle.2026.106563
Received 19 February 2026; Received in revised form 2 July 2026; Accepted 7 August 2026
Available online 19 August 2026
0167-6911/© 2026 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

slow system variations without interfering with the system dynamics. This separation of time scales, however, limits the choice of the dither frequency. More specifically, the dither frequency should be chosen to ensure sufficient excitation for gradient estimation (satisfactory high signal-to-noise ratio (SNR)) while not interfering with system dynamics [14].

Increasing the dither frequency without consideration of system dynamics can lead to stationary solutions that deviate from the true optimum, and multiple solutions may exist [15]. This behavior is exposed in [16], which demonstrated that the error between the true optimum and the steady state to which the system converges is proportional to the square of the excitation amplitude and frequency. The same work also showed, through linear analysis of the system under optimization, that a zero moves from the left half-plane to the right half-plane when crossing the optimum. This observation was further examined in [17], where this behavior was shown to be strongly dependent on the definition of the objective function, and may therefore not always be present.

Several advanced approaches have been proposed to provide a (partial) solution to the time-scale separation limitations in classical extremum seeking control. In various publications, a proportional term is added to the classical ESC scheme in addition to the integral term to improve convergence behavior [18,19]. Furthermore, a notable research line is phasor-based ESC, which estimates the amplitude and phase (phasor) of the system's response to the periodic dither perturbation, rather than estimating the objective gradient as in the classical demodulation strategy [20]. The advantage of this method is its robustness to both large and variable system phase lags. The phasor-estimation scheme is further extended by adaptively adjusting the perturbation amplitude in the neighborhood of the optimum, thereby improving convergence robustness and accuracy [21]. An alternative approach to addressing time-scale separation limitations is to use frequency-domain techniques, such as recursive Fourier analysis and online system identification, to estimate the local frequency response. This approach enables gradient estimation via complex curve fitting, allowing accurate optimization even when dither frequencies overlap with plant dynamics [22]. Finally, a semi-global stability result for Wiener–Hammerstein plants using high dither frequencies is presented in [23], leading to fast convergence in an arbitrarily small region of the optimum.

To enable the analysis of the time-scale separation limitations of classical ESC, a formal Wiener system analysis framework is provided in [24,25]. It is observed that ESC relies on *good phase difference detection* between the input and the output. In specific dynamical systems, such phase difference properties occur only at very low frequencies. Consequently, low dither frequencies lead to prohibitively slow convergence, rendering classical ESC impractical for certain applications. To overcome this challenge, a pre-compensated ESC scheme uses a phase-locked loop (PLL) to retrieve gain and phase (i.e., phasor) information at the excitation frequency, thereby improving gradient estimation. While the presented method demonstrably accelerates convergence, the authors focus on convergence speed, and no analysis is provided regarding the consistency of the steady-state solution with respect to the true optimum.

This paper investigates the convergence properties of ESC when applied to dynamical systems with particular emphasis on how the formulation of the performance objective influences the convergence behavior. Two types of objectives are considered: state-dependent and bilinear (state-input dependent) objectives, using a Wiener-type model structure to capture static output nonlinearities that follow linear dynamics. To analyze the convergence characteristics, two complementary analysis methods are employed: (i) sequential linearization combined with frequency-domain examination [17] and (ii) a more formal signal-based analysis. These methods provide both intuitive insights and explanations, facilitating the development of enhanced ESC methods that improve convergence speed and consistency with the true optimum. The contributions of this paper are:

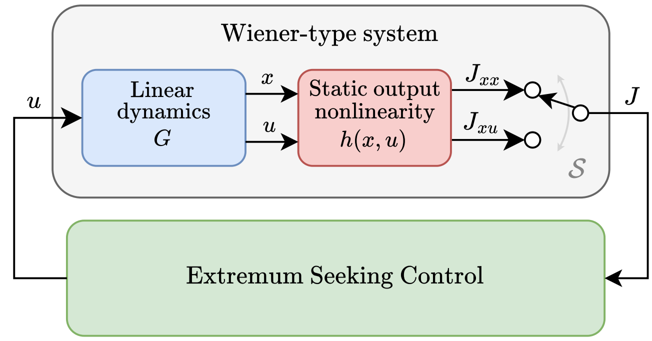

Fig. 1. Structure of the dynamical Wiener-type system $\mathcal{S}$, with linear-time invariant dynamics G followed by a static nonlinearity h. The scalar input u yields the state x and a direct-feedthrough term as outputs. Both are used to compute the performance objectives J_{xx} and J_{xu}. The static input–output mapping between u and the objectives has equal maximizers. The convergence properties of extremum seeking control based on these objectives are of interest.

1. Formulating a Wiener-type model framework for analyzing ESC convergence properties in dynamical systems with state-dependent and bilinear performance objectives.
2. Demonstrating that ESC convergence behavior can differ significantly for performance objectives with identical steady-state optima, particularly when bilinear objectives are used in dynamical systems.
3. Analyzing the convergence phenomena using sequential linearization and frequency-domain examinations, and signal-based analyses.
4. Proposing enhanced ESC methods enabling faster and consistent convergence by overcoming time-scale separation constraints, allowing for higher dither frequencies. The efficacy of the proposed methods is demonstrated in simulations.

The paper is organized as follows. Sections 2 and 3 formalize the model structure for system analysis and revisit the theory of dither-demodulation-based extremum seeking control. Section 4 exposes the convergence phenomenon by contrasting a steady-state analysis with an illustrative example optimizing a dynamical system that reveals *unexpected* ESC convergence behavior. Section 5 provides detailed analyses of this behavior using frequency-domain methods and signal-based analyses. Based on the presented analyses, Section 6 proposes methods to improve ESC convergence properties, and showcases their efficacy in simulations. Finally, conclusions are drawn in Section 7.

Notation

Throughout the paper, scalars and vectors are denoted by lowercase letters, and matrices by uppercase letters. Operating points and mean values are indicated by an overbar $\bar{(\cdot)}$, and time-derivatives are denoted by $\dot{(\cdot)}$. The notation $\hat{(\cdot)}$ implies estimated quantities and optimal values are indicated by $(\cdot)^*$.

2. System and model structure

This section defines the system and its model structure subject to the ESC scheme. Fig. 1 shows the closed-loop dynamical system under optimization of ESC. This work considers systems with a Wiener-type model structure denoted by $\mathcal{S}$, where, for analysis, the system model is assumed to be known. This model type is chosen for its ability to capture static nonlinearities following linear dynamics, which aligns with the structure of many practical systems optimized using the ESC scheme.

The goal of ESC is to drive the input u toward the optimizer u^* that maximizes the objective J. Maximization is considered here, but the results apply to minimization by sign reversal. The system, of which the input is subject to the ESC algorithm, is defined by the state space model:

$$\dot{x} = f(x,u) = \frac{1}{\tau}\left(-\gamma(x) + c\alpha(u)\right), \tag{1}$$

and

$$J = h(x,u) = \begin{cases} J_{xx} = -\gamma(x)x, & \text{(2a)} \\ J_{xu} \;= -c\alpha(u)x, & \text{(2b)} \end{cases}$$

with $f : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^n$ being the state equation, $\gamma : \mathbb{R}^n \to \mathbb{R}^n$ and $\alpha : \mathbb{R}^p \to \mathbb{R}^n$ the state and input functions, $h : \mathbb{R}^n \times \mathbb{R}^p \to \mathbb{R}^q$ the output function, with state $x \in \mathbb{R}^n$ and input $u \in \mathbb{R}^p$. The system order is n, and the number of inputs p is set to 1; this assumption is made because SISO ESC is considered, and to facilitate a concise analysis. The time constant $\tau \in \mathbb{R}$ defines the system's dynamic response speed, while $c \in \mathbb{R}$ is an input scaling coefficient.

By the Wiener-type model structure, the system consists of a linear time-invariant (LTI) block G representing (1), followed by a static output nonlinearity $h(\cdot,\cdot)$. The input u drives system G, whose state x is assumed to be available through measurement or estimation. This state, along with a possible direct feedthrough of u, is passed to $h(x,u)$, which generates the performance objective(s). Given the scope of this paper, restricted to Wiener-class systems, the functions γ, α are to be understood as linear (possibly affine) functions. Thus, the performance objectives in (2) are quadratic, for the case of J_{xx}, or bilinear, J_{xu}. For the latter, bilinear refers to objectives in which the nonlinearity arises from the product of the system state and control input, and is characterized by linear dependence on each variable when the other is held fixed. As shown in Fig. 1, either the state-dependent objective J_{xx} or the bilinear (state-input dependent) objective J_{xu} is selected as the performance objective output J for maximization by the ESC algorithm. Respectively, the state function $\gamma(x)$ in J_{xx} depends solely on the state, and the input function $\alpha(u)$ in J_{xu} solely on the input. This work focuses on single-objective ESC schemes optimizing a single input variable.

While the modeling framework introduced in this section allows for general functional forms of $\gamma(x)$ and $\alpha(u)$, the convergence analysis presented in this paper explicitly focuses on affine definitions of these functions. This modeling choice enables transparent and tractable analyses of extremum seeking control convergence properties, allowing the fundamental mechanisms associated with state-dependent and bilinear objectives to be exposed. Although more general smooth nonlinear mappings may be accommodated locally through linearization around operating points, all analytical results presented in this paper rely on affine $\gamma(x)$ and $\alpha(u)$. A systematic investigation of fully nonlinear mappings lies beyond the scope of the present work and constitutes an interesting direction for future research.

Throughout this paper, the term *Wiener-type model* is used to denote a generalized version of the classical Wiener model structure. In its strict definition, a Wiener model consists of linear time-invariant (LTI) dynamics followed by a static nonlinearity. The formulation adopted in this work allows for generally nonlinear mappings in the state equation through the functions $\gamma(x)$ and $\alpha(u)$, while the analysis relies on their affine approximations obtained via local linearization around steady-state operating points. This generalized Wiener-type representation is employed to facilitate local linearization and frequency-domain analysis of extremum seeking control convergence properties, which are the primary focus of the present study.

3. Extremum seeking control

Extremum seeking control is a model-free optimization technique that optimizes quasi-static inputs to drive systems toward the optimum of their steady-state input–output mapping. This section outlines the classical dither-demodulation ESC scheme used throughout this work [5], and adopts the notation in [26]. Fig. 2 illustrates its structure. The goal of the algorithm is to find a (local) optimum of the objective J by estimating its gradient with respect to the control input u.

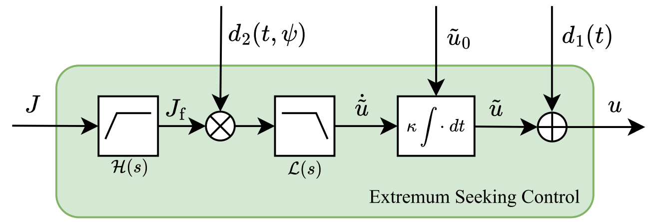

Fig. 2. Classical implementation of a dither-demodulation extremum seeking control scheme. The objective is fed to the ESC subsystem and is subsequently high-pass filtered, demodulated, and low-pass filtered to estimate the objective function gradient with respect to the system input. This gradient is numerically integrated until convergence. The dither excitation signal is added in the end.

To achieve this estimate, the input signal u is composed of a quasi-steady-state term $\tilde{u}$, which is perturbed by an appropriately defined periodic dither excitation signal. The objective of the ESC algorithm is for the time derivative of the quasi-steady-state input part to approximate

$$\dot{\tilde{u}} \simeq \kappa \frac{\partial J}{\partial u}, \tag{3}$$

where $\dot{\tilde{u}}$ is the time derivative of the quasi-steady-state control input under optimization, and $\kappa \in \mathbb{R}$ is the integral (or convergence) gain of the algorithm.

The ESC scheme estimates the gradient $\partial J/\partial u$ via periodic excitation, filtering, and demodulation operations, as illustrated in Fig. 2. The measured performance objective J is high-pass filtered by $\mathcal{H}(\omega_{\mathcal{H}})$ to remove the DC-offset, resulting in J_{f}, with s the Laplace operator and $\omega_{\mathcal{H}}$ the filter's cut-in frequency. The resulting zero-mean signal is subsequently demodulated by multiplication with the signal

$$d_2(t,\psi) = \sin(\omega_{\mathrm{d}} t + \psi), \tag{4}$$

with unity amplitude and dither frequency ω_{d}, the demodulation signal. The term ψ is chosen for phase compensation, accounting for the combined phase loss due to the system dynamics and the high-pass filter, and improving the gradient estimation's accuracy to enhance convergence characteristics [27]. Under appropriate design, the result of this demodulation operation is a signal whose steady-state contribution contains gradient information, and a periodic contribution at twice the dither frequency.

The subsequent low-pass filtering operation by $\mathcal{L}(\omega_{\mathcal{L}})$ attenuates noise and the higher harmonic dither frequency, and $\omega_{\mathcal{L}}$ is the filter's cut-off frequency. This results in the signal $\dot{\tilde{u}}$ of which the steady-state component represents the desired gradient $\partial J/\partial u$. Then, the signal $\dot{\tilde{u}}$ is numerically integrated over time and scaled by κ as a scheme tuning variable balancing convergence speed and stability, resulting in the quasi-steady-state control input $\tilde{u}$. Finally, the dither signal

$$d_1(t) = A_{\mathrm{d}} \sin(\omega_{\mathrm{d}} t), \tag{5}$$

with amplitude A_{d} is added to $\tilde{u}$ and forms the combined input

$$u = d_1 + \tilde{u}, \tag{6}$$

to the system under optimization. The dither amplitude A_{d} is calibrated as a trade-off between signal-to-noise ratio (SNR) and permissible system excitation levels. In classical ESC implementations, the dither frequency should be chosen on the slow-to-medium time scale to ensure desirable convergence properties – an aspect particularly relevant in this work.

The respective high-pass filter and low-pass filter, cut-in and cut-off frequencies are related to the dither frequency by the empirical relations

$$\omega_{\mathcal{H}} = \omega_{\mathrm{d}}/5, \quad \omega_{\mathcal{L}} = 2\omega_{\mathrm{d}}. \tag{7}$$

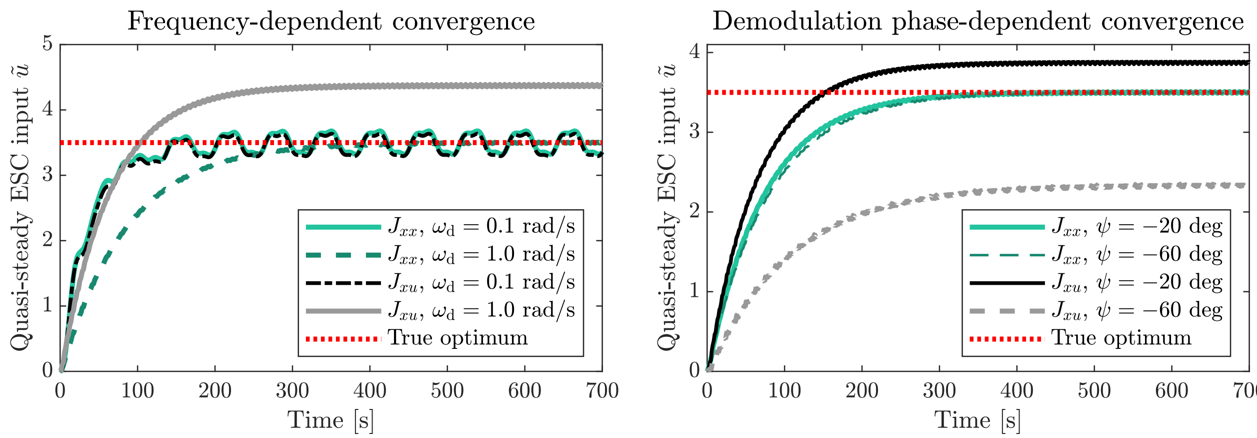

(a) Dither frequencies were selected in the slow-medium and fast time scales for both optimization objectives (demodulation phase constant $\psi = 0$ deg). While both objectives converge to the true optimum at the slower time scale, the bilinear objective exhibits a steady-state offset from the optimum at the faster time scale. Results indicate time-scale sensitivity and convergence consistency; furthermore, lower dither frequencies lead to larger residual oscillations around the optimum.

(b) Various demodulation phase offsets were employed for both optimization objectives at a constant fast-time-scale dither frequency ($\omega_d = 1.0$ rad/s). For the bilinear objective, results show inconsistent convergence to the true optimum across the phase offsets. In contrast, for the state-dependent objective, phase variation affects ESC convergence speed and stability [7], while maintaining consistent convergence.

Fig. 3. Extremum seeking control optimization trajectories based on the dynamical system objectives J_{xx} and J_{xu}, with the true steady-state optimum indicated by the red dashed line. Two convergence analyses are presented: (a) choosing the dither frequency in the slow-mid and fast time scales, and (b) varying the demodulation phase in the fast time scale.

The described real-time process proceeds until the gradient diminishes and the control input converges to a neighborhood of the optimal control input $\tilde{u}^*$. More specifically, the ESC algorithm converges to a neighborhood of the optimal input [6], whenever the following three conditions are met

$$\frac{\partial J}{\partial u}(u^*) = 0 \text{ (optimum exists)}, \tag{8a}$$

$$\frac{\partial^2 J}{\partial u^2}(u^*) < 0 \text{ (maximum)}, \tag{8b}$$

$$\frac{\partial J}{\partial u}(u^* + \zeta)\zeta < 0 \ \forall \ \zeta \in Z, \tag{8c}$$

where (8c) indicates that u^* is a local maximizer in some region Z. The conditions given in (8) are steady-state conditions; however, in dynamical systems, convergence properties are susceptible to input–output dynamics and the frequency of the dither signal. The next section formalizes and illustrates the problem addressed in this paper.

4. Problem formalization

In Section 2, two objectives were introduced. This section formalizes the problem addressed in this paper. First, a steady-state analysis confirms that both objectives (state-dependent and bilinear) lead to the same optimal input. Then, ESC is applied to a dynamical system with such objectives that share identical static input–output characteristics, to evaluate how each objective affects the algorithm's ability to converge to the true optimum.

4.1. Steady-state analysis

To show that the two functions in (2) have the same steady-state optimizer, an example state function $\gamma(x) = x + r$ is taken with $r \in \mathbb{R}$ a constant scalar state offset, and the input function is taken as $\alpha(u) = u$. Evaluating the objectives in (2) results in

$$\begin{aligned} J_{xx} = -\gamma(x)x &= -(x + r)(cu - \tau\dot{x} - r) \\ &= (\tau\dot{x} - cu)(cu - \tau\dot{x} - r), \end{aligned} \tag{9}$$

$$J_{xu} = -c\alpha(u)x = -cu(cu - \tau\dot{x} - r). \tag{10}$$

For clarity, it is noted that the dynamical expressions containing the term $\dot{x}$ are used in the presented derivation to show and maintain a direct link with the underlying continuous-time system dynamics. The resulting steady-state expressions are obtained by explicitly imposing $\dot{x} = 0$.

Assuming steady-state conditions (i.e., $\dot{x} = 0$), it becomes apparent that both objectives are equal. This conclusion can also be directly drawn from (1) by taking the partial derivative with respect to the perturbing input:

$$\left.\frac{\partial J_{xx}}{\partial u}\right|_{\dot{x}=0} = \left.\frac{\partial J_{xu}}{\partial u}\right|_{\dot{x}=0} = cr - 2c^2u = 0 \ \rightarrow \ u^* = \frac{r}{2c}. \tag{11}$$

The next section shows how implementing the ESC algorithm may have unintended consequences on convergence when the system dynamics interfere at the dither frequency.

4.2. Dynamical analysis using ESC

In this section, equal steady-state properties are now complemented with system dynamics, and extremum seeking control is employed to find the optimizer u^*_{xx} and u^*_{xu} for both the state-dependent and bilinear objectives, respectively, with results shown in Fig. 3. The system parameters are chosen as $c = 1$, $r = 7$, and $\tau = 1$, and according to the result in (11), the steady-state optimum is expected to be at $u^*_{xx} = u^*_{xu} = 3.5$. The ESC algorithm is configured with a dither amplitude of $A_d = 4.0$ and a convergence gain of $\kappa = 0.01$.

To demonstrate the effects of time-scale separation on the convergence result using both objectives, two convergence analysis cases are presented. First, the dither frequency for the ESC algorithm is chosen as $\omega_d = 0.1$ rad/s and $\omega_d = 1.0$ rad/s. The former is in the slow-medium time scale, whereas the latter is in the fast time scale – the latter is deliberately chosen to violate the time-scale separation principle. The low-pass and high-pass filter parameterizations are selected according to (7). The second analysis is performed by varying the demodulation phase in the fast time scale. Demodulation phases of $\psi = -20$ deg and $\psi = -60$ deg are chosen.

Fig. 3 presents ESC convergence results for both objectives subject to distinct dither frequencies. The red dotted line indicates the true steady-state optimum. Fig. 3(a) illustrates how varying the dither frequency across the slow-medium and fast time scales affects convergence. The state-dependent objective consistently converges near the true optimum, whereas the bilinear objective exhibits an offset to the true optimum when the dither frequency is in the system dynamic frequency range. Additionally, both objectives demonstrate that lower dither frequencies lead to increased residual oscillations around the optimum [16].

Fig. 3(b) shows that for a constant fast dither frequency of $\omega_{\mathrm{d}} = 1.0$ rad/s, the convergence inconsistency with respect to the true optimum in the bilinear objective also varies with the demodulation phase. For the state-dependent objective, the ESC algorithm converges consistently to similar neighborhoods around the true optimum. In this case, the demodulation phase influences the convergence rate by enhancing the correlation between the demodulation signal and the objective signal's frequency content [28].

As both objectives share the same steady-state optimum, the observed discrepancy – particularly in the bilinear case – defines the core problem addressed in this paper and motivates subsequent analysis and proposed resolution. While this study focuses on single-variable, single-objective ESC, the analytical and solution frameworks introduced here are readily extendable to multivariable systems according to the definitions in [7].

5. ESC convergence analysis

This section presents a series of analyses that explain the convergence behavior of ESC for the considered state-dependent and bilinear objectives and introduces shared theoretical preliminaries. Throughout this section, the input scaling coefficient is taken as unity ($c = 1$). This normalization is adopted in this section for analytical convenience and to simplify the expressions used to study the qualitative convergence mechanisms.

Two analyses offer distinct perspectives on the problem under consideration. First, a frequency-domain investigation based on successive linearizations provides intuitive insights into the system's optimization behavior under ESC. Building on these insights, a more formal signal-based derivation is presented to deepen the understanding of the observed convergence characteristics.

5.1. Preliminaries

This section outlines the theoretical foundations and notation used throughout this paper. It begins with a set of relevant trigonometric identities:

$$\sin^2(x) = \frac{1}{2}(1 - \cos(2x)), \tag{12a}$$

$$\sin(x)\sin(x+y) = \cos(y)\sin^2(x) + \sin(y)\sin(x)\cos(x), \tag{12b}$$

$$\sin(x)\cos(x) = \frac{1}{2}\sin(2x), \tag{12c}$$

$$\sin(x) + \sin(x+y) = \sin(x) + \cos(y)\sin(x) + \cos(x)\sin(y), \tag{12d}$$

$$\sin(x+y) + \sin(x+z) = \frac{1}{2}\left(\cos(y-z) - \cos(2x+y+z)\right). \tag{12e}$$

Next, consider the linear dynamic system G, which forms part of the Wiener-type system defined in (1), subject to a periodic input signal defined in (6). The resulting steady-state response is given by:

$$x = M\sin(\omega_{\mathrm{d}}t + \Psi) + M_0\tilde{u} - r, \tag{13}$$

and taking the time derivative yields:

$$\dot{x} = \omega_{\mathrm{d}} M\cos(\omega_{\mathrm{d}}t + \Psi), \tag{14}$$

where the quasi-static component $\tilde{u}$ is treated as time-invariant. The steady-state response consists of (i) a periodic component at the dither frequency, (ii) a constant contribution due to the quasi-steady control input $\tilde{u}$, and (iii) a constant offset induced by the chosen state mapping. This decomposition follows directly from the linearity of the dynamical system driven by the composite dither input defined earlier in (6). By linearity and superposition, the state response can be written as the sum of the corresponding individual responses. The first term represents the periodic steady-state response to the sinusoidal dither excitation ω_{d}, with magnitude $M = A_{\mathrm{d}}|G(\mathrm{j}\omega_{\mathrm{d}})|$ and phase shift $\Psi = \angle G(\mathrm{j}\omega_{\mathrm{d}})$ of the system frequency response, with j the imaginary unit. The second term arises from the quasi-steady-state control input $\tilde{u}$ and is scaled by the DC gain $M_0 = |G(j0)|$ of the linear dynamical system. The third constant offset term $-r$ originates from the chosen state mapping, which deliberately shifts the equilibrium of the state dynamics.

For analysis purposes, theoretically ideal filters with perfect frequency-selective properties indicated by $(\cdot)^*$ are defined [29]. An ideal low-pass filter is expressed in the frequency domain as:

$$\mathcal{L}^*(\omega_{\mathcal{L}}) = \begin{cases} 1, & |\omega| \le \omega_{\mathcal{L}} \\ 0, & \text{otherwise,} \end{cases} \tag{15}$$

and the corresponding ideal high-pass filter is:

$$\mathcal{H}^*(\omega_{\mathcal{H}}) = \begin{cases} 1, & |\omega| > \omega_{\mathcal{H}} \\ 0, & \text{otherwise.} \end{cases} \tag{16}$$

These filters can be combined to form an ideal band-pass filter by appropriately selecting the lower and upper cut-off frequencies such that $\omega_{\mathcal{B},\mathrm{U}} > \omega_{\mathcal{B},\mathrm{L}}$:

$$\mathcal{B}^*(\omega_{\mathcal{B},\mathrm{L}}, \omega_{\mathcal{B},\mathrm{U}}) = \mathcal{L}^*(\omega_{\mathcal{B},\mathrm{U}})\mathcal{H}^*(\omega_{\mathcal{B},\mathrm{L}}), \tag{17}$$

and in a similar manner, an ideal notch filter (or band-stop filter) can be constructed as the complement of the band-pass filter, rejecting frequencies within a specified band while passing all others:

$$\begin{aligned}\mathcal{N}^*(\omega_{\mathcal{N},\mathrm{L}}, \omega_{\mathcal{N},\mathrm{U}}) &= 1 - \mathcal{B}^*(\omega_{\mathcal{N},\mathrm{L}}, \omega_{\mathcal{N},\mathrm{U}}) \\ &= \mathcal{L}^*(\omega_{\mathcal{N},\mathrm{L}}) + \mathcal{H}^*(\omega_{\mathcal{N},\mathrm{U}}),\end{aligned} \tag{18}$$

The band-pass filter passes frequencies $\omega_{\cdot,\mathrm{L}} < |\omega| \le \omega_{\cdot,\mathrm{U}}$, while the notch filter rejects this frequency band. The center frequencies of both filters are defined as

$$\omega_{\mathcal{B}} = \omega_{\mathcal{N}} = (\omega_{\cdot,\mathrm{L}} + \omega_{\cdot,\mathrm{U}})/2. \tag{19}$$

5.2. Frequency-domain analysis

This section presents a frequency-domain analysis of the dynamical system subject to the ESC scheme. To this end, successive linearizations of the nonlinear Wiener-type system $\mathcal{S}$ are derived at steady-state operating points around the optimal input for both objective outputs. Each linearized model is therefore valid locally around its corresponding operating point. By analyzing successive linearizations for ESC-driven inputs moving across the steady-state optimum, classical frequency-domain control tools can be used to examine how the system dynamics, ESC scheme, and objective function definition jointly influence convergence behavior near the optimum. This approach provides local mechanistic insights into ESC operation in the presence of system dynamics, rather than a global system characterization.

The state and output equations of $\mathcal{S}$ in (1) and (2) are analytically linearized, resulting in the linear state-space system:

$$A = \left.\frac{\partial f(x,u)}{\partial x}\right|_{\bar{x},\bar{u}} = -\frac{1}{\tau}\frac{\partial \gamma}{\partial x}, \tag{20a}$$

$$B = \left.\frac{\partial f(x,u)}{\partial u}\right|_{\bar{x},\bar{u}} = \frac{1}{\tau}\frac{\partial \alpha}{\partial u}, \tag{20b}$$

$$C = \left.\frac{\partial h(x,u)}{\partial x}\right|_{\bar{x},\bar{u}} = \left\{ \begin{array}{l} -\frac{\partial \gamma(\bar{x})}{\partial x}\bar{x} - \gamma(\bar{x}) \\ -\alpha(\bar{u}) \end{array} \right. , \tag{20c}$$

$$D = \left.\frac{\partial h(x,u)}{\partial u}\right|_{\bar{x},\bar{u}} = \left\{ \begin{array}{l} 0 \\ -\frac{\partial \alpha(\bar{u})}{\partial u}\bar{x} \end{array} \right. , \tag{20d}$$

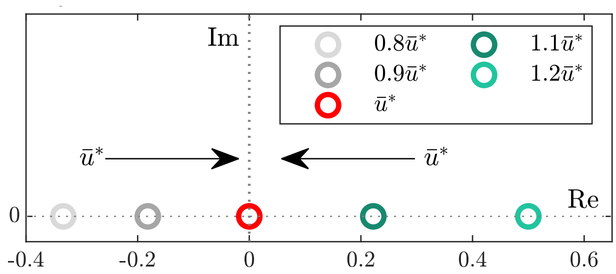

Fig. 4. For the bilinear objective, the system exhibits alternating minimum-phase and nonminimum-phase behavior when the input crosses the optimal value. More specifically, a zero crosses the imaginary axis at $\alpha(\bar{u}) = -(\partial\gamma/\partial x)\bar{x}$ for the considered system, shifting between the left half-plane and the right half-plane.

where $A \in \mathbb{R}^{n\times n}$ is the state matrix, $B \in \mathbb{R}^{n\times p}$ the input matrix, $C \in \mathbb{R}^{q\times n}$ the output matrix, and $D \in \mathbb{R}^{q\times p}$ the direct-feedthrough matrix. The input scaling coefficient is normalized to $c = 1$. The operating points $\bar{x}, \bar{u}$ are computed using the analytic steady-state solution given in Section 4.1 for values above and below the optimum.

Using $C(sI - A)^{-1}B + D$, unique transfer function realizations of the state-space system are computed:

$$\mathbf{H} = \begin{bmatrix} H_{xx} \\ H_{xu} \end{bmatrix} = \begin{bmatrix} N_{xx}/\mathcal{D} \\ N_{xu}/\mathcal{D} \end{bmatrix}, \tag{21}$$

with

$$N_{xx} = -\frac{\partial\alpha}{\partial u}\left(\gamma + \frac{\partial\gamma}{\partial x}\bar{x}\right), \quad N_{xu} = -\frac{\partial\alpha}{\partial u}\left(\alpha + \tau\bar{x}s + \frac{\partial\gamma}{\partial x}\bar{x}\right),$$
$$\mathcal{D} = \tau s + \frac{\partial\gamma}{\partial x},$$

where s is the Laplace variable. Note that the transfer function denominator (and thus the system poles) is identical for both objectives. However, the linear state-space system for the bilinear objective J_{xu} includes a nonzero direct-feedthrough term (see (20d)), resulting in a varying zero location depending on the linearization point, as shown in Fig. 4. Its implications will be discussed shortly.

For the frequency-domain analysis, the same system parameterization from Section 4 is used. System frequency responses are evaluated near the steady-state operating point corresponding to the optimal input $\bar{u}^*$ over the frequency range $\Omega = \{\omega \in \mathbb{R} \mid 10^{-3} \leq \omega \leq 10^{2}\}$. Fig. 5 presents the results, with the dither frequency from Section 4.2 marked by vertical dotted lines.

Fig. 5(a) shows the frequency responses of the state-dependent objective $H_{xx}(\mathrm{j}\omega)$, revealing two key characteristics that favor ESC convergence. First, the magnitude of the response diminishes as the input approaches the optimum, effectively slowing down the convergence rate near the optimum. Second, the phase response exhibits a shift of 180 deg across the entire frequency range when the input crosses $\bar{u}^*$, indicating a sign change in the system response phase. This combined behavior ensures a reliable convergence mechanism.

In contrast, Fig. 5(b) presents the frequency responses for the bilinear objective $H_{xu}(\mathrm{j}\omega)$, showing more complex behavior across the frequency spectrum. At lower frequencies, in the region $10^{-3} \leq \omega \leq 10^{-2}$ rad/s, the system exhibits beneficial magnitude and phase behavior similar to H_{xx}, enabling convergence when the dither frequency is chosen within this range. However, in the higher frequency region $10^{\circ} \leq \omega \leq 10^{2}$ rad/s, these favorable properties are absent. The previously described beneficial magnitude and phase properties no longer hold, complicating ESC convergence and resulting in inconsistent convergence to the true optimizer. A more in-depth discussion with application to wind turbine torque controller optimization is provided in [17].

This deterioration at higher frequencies can be attributed to a zero in the bilinear transfer function, which is absent in the state-dependent transfer function. As shown in Fig. 4, these effects are induced by a varying zero-location that moves from the left-half plane (LHP) to the right-half plane (RHP) when crossing the optimum input. As a result, the system H_{xu} alternates between a minimum-phase and nonminimum-phase system for equilibrium points below and above $\bar{u}^*$, introducing nontrivial phase distortions that hinder ESC performance. The analytical expression for when the zero is located at the origin, that is, on both the real (Re) and imaginary (Im) axis, is found by finding the zeros of the numerator N_{xu}. For the considered system, this results in $\alpha(\bar{u}) = -(\partial\gamma/\partial x)\,\bar{x}$. This result has been described earlier by [16], and is further demonstrated in [17].

This qualitative frequency-domain analysis of state-dependent and bilinear objectives reveals convergence differences between the two. Convergence of bilinear objectives appears to be more challenging, even when their steady-state optimum coincides with the corresponding state-dependent objective. For classical ESC schemes, the presented analysis for the bilinear objective motivates time-scale separation and the careful selection of low dither frequencies outside the fast time scale to ensure reliable convergence. Notably, state-dependent objectives exhibit inherently advantageous characteristics, such as diminishing gain and consistent phase difference behavior (sign change) near the optimum. These properties support convergence even at higher dither frequencies, thereby relaxing the need for strict time-scale separation. The next section provides a signal-based analysis to formally substantiate these observations.

5.3. Signal-based analysis

This section presents an open-loop, signal-based analysis of the combined system $\mathcal{S}$ in series with the ESC algorithm (without integrator). Fig. 6 shows the schematic block diagram of this open-loop configuration, mapping the system input u to the ESC-estimated cost function gradient $\hat{u}$. Separate analyses are carried out for the two objective outputs $J = \{J_{xx}, J_{xu}\}$. The aim is to analytically derive the steady-state convergence value for both objectives, accounting for the dynamical properties, filtering, and signal operations in the open-loop configuration.

The ideal filters $\mathcal{H}^*$ and $\mathcal{L}^*$ used in this analysis serve as analytical tools that enable ideal separation of frequency components in the signal-based derivations. Their use allows for the derivation of closed-form expressions for steady-state optimizing inputs. Since the analysis targets the steady-state input solution at zero frequency (DC), the resulting expressions remain valid for practical ESC implementations using implementable filters, provided they are tuned to approximate the frequency-selective properties in the analysis. In practice, non-ideal filtering may affect transient behavior or lead to small steady-state oscillations around the predicted optimum, for instance, due to aggressive ESC convergence gain tuning, but does not introduce a bias in the steady-state optimizing input under standard ESC assumptions.

The state and input functions are defined as in the previous section, with normalized input scaling: $\gamma(x) = x + r$ and $\alpha(u) = u$, and with identical parameter values. A step-by-step analysis is provided for both the state-dependent and bilinear objectives.

State-Dependent Objective J_{xx}: The ideal state-dependent objective function (from (2a)) is

$$J_{xx} = -\gamma(x)x = -(x + r)\,x. \tag{22}$$

Substituting the prerequisites from Section 5.1 gives

$$J_{xx} = -(M\sin(\omega_{\mathrm{d}}t + \Psi) + M_0\tilde{u}) \cdot (M\sin(\omega_{\mathrm{d}}t + \Psi) + M_0\tilde{u} - r), \tag{23}$$

which expands to

$$J_{xx} = -M^2\sin^2(\omega_{\mathrm{d}}t + \Psi) - M_0^2\tilde{u}^2 + M_0\tilde{u}r - (2MM_0\tilde{u} - Mr)\sin(\omega_{\mathrm{d}}t + \Psi). \tag{24}$$

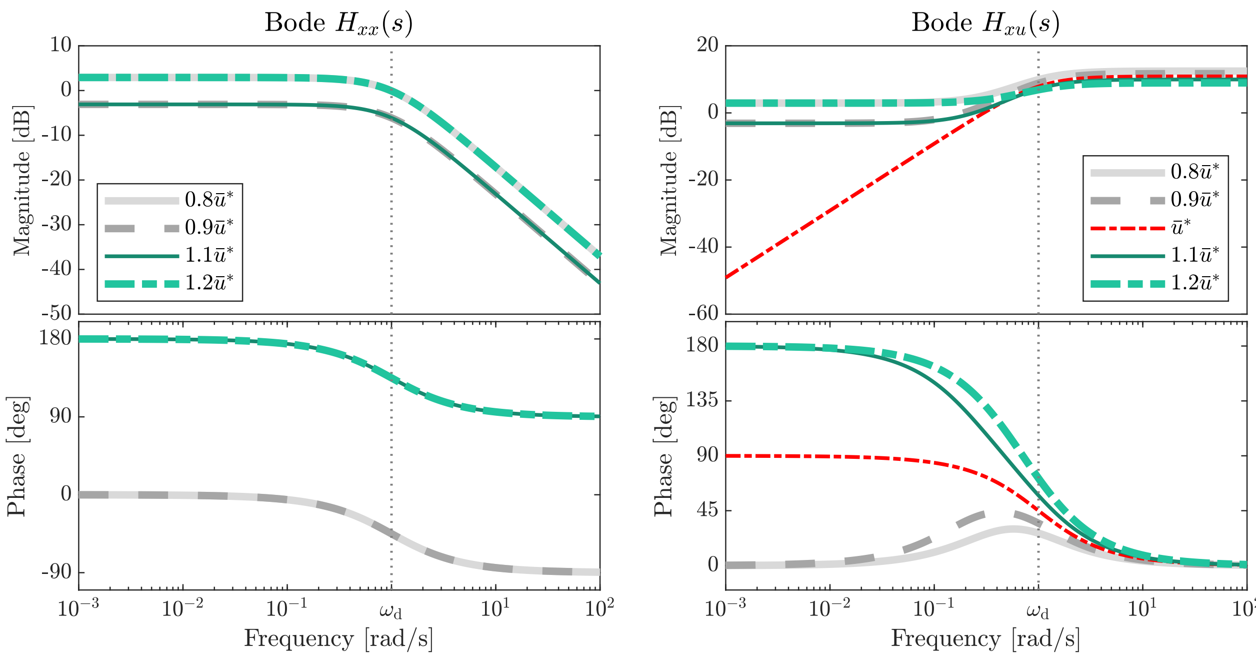

(a) State-dependent objective J_{xx} (b) Bilinear objective J_{xu}

Fig. 5. Frequency responses of linearized models H_{xx} and H_{xu} near the optimal operating point input $\bar{u}^*$. For J_{xx}, favorable ESC convergence behavior is facilitated across all frequencies by both a diminishing magnitude transfer and a 180 deg phase shift when the input crosses $\bar{u}^*$. The frequency response for the operating point at $\bar{u}^*$ is not shown, as the transfer function equals zero at the optimum. In contrast, for J_{xu}, the described properties appear only at low frequencies, emphasizing and explaining the importance of time-scale separation.

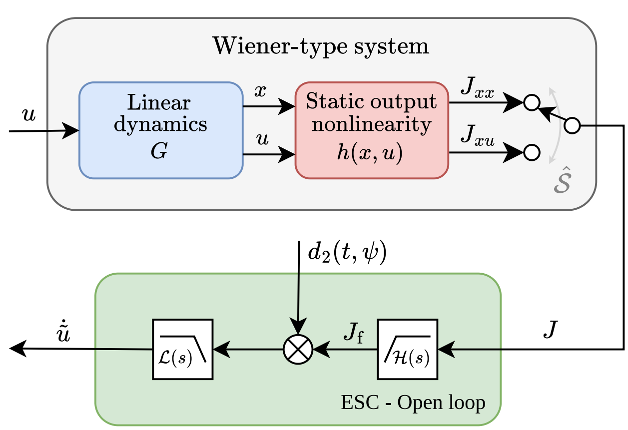

Fig. 6. Open-loop Wiener-type system and ESC scheme without integrator. This configuration facilitates steady-state convergence analysis by tracking signals from system input u to ESC gradient estimate $\hat{\dot{u}}$.

Applying trigonometric identities in (12) yields

$$\begin{aligned} J_{xx} = &\frac{M^2}{2}\cos(2\omega_d t + 2\Psi) \\ &+ (Mr - 2M M_0 \tilde{u})\sin(\omega_d t + \Psi) \\ &- \frac{M^2}{2} - M_0^2\tilde{u}^2 + M_0\tilde{u}r. \end{aligned} \tag{25}$$

The term in the first line is at two times the dither frequency, those in the second line are at the dither frequency, and the last line contains constants. The signal is filtered with a high-pass filter $\mathcal{H}(\omega_{\mathcal{H}} < \omega_d)$ to remove the constant, and introduces a phase θ at ω_d. The subsequent demodulation with $d_2(t,\psi)$ shifts the frequency content by $(\omega \pm \omega_d)$. After demodulation, only the steady-state (DC) component of $\hat{\dot{u}}$ is used for gradient estimation. Therefore, only terms at the dither frequency before demodulation contribute to a steady-state contribution after demodulation. The signal before demodulation and after high-pass filtering is represented by

$$J_{xx}^{f} = (Mr - 2M M_0\tilde{u})\sin(\omega_d t + \Psi + \theta) + \mathcal{O}(>\omega_d), \tag{26}$$

where $\mathcal{O}(>\omega)$ denotes higher-frequency components whose exact form is not relevant for the steady-state analysis. Next, the signal J_{xx}^{f} is demodulated by multiplication with the unity-amplitude demodulation signal $d_2(t,\psi)$, resulting in:

$$\hat{\dot{u}} = \frac{M}{2}\left(r - 2M_0\tilde{u}\right)\cos\left(\psi - (\Psi + \theta)\right) + \mathcal{O}(>\omega_d) \tag{27}$$

and a low-pass filtering operation by $\mathcal{L}^*$ is performed to eliminate periodic components, leading to elimination of the $\mathcal{O}(>\omega_d)$ term, and further simplification results in:

$$\hat{\dot{u}} = \frac{M}{2}(r - 2M_0\tilde{u})\cos(\psi - (\Psi + \theta)). \tag{28}$$

To determine the ESC convergence point, the expression is equated to zero, resulting in the optimal input to the system

$$\boxed{u_{xx}^* = \frac{r}{2M_0}}. \tag{29}$$

Thus, for a state-dependent objective function, the ESC final steady-state convergence result is only dependent on the steady-state system gain and the state offset r in $\gamma(x)$, and is independent of dynamic system properties. This agrees with the frequency-domain analysis, and the steady-state optimal result in (11). As becomes clear from (27), and in contrast to the final convergence value, the convergence dynamics are influenced by dynamic system properties. In particular, the demodulation introduces a multiplicative factor $\cos(\psi - (\Psi + \theta))$, which affects both convergence rate and gradient sign, and may suppress adaptation when it vanishes. The steady-state convergence point remains unaffected, provided that $\cos(\psi - (\Psi + \theta)) \neq 0$. The following section presents a similar analysis for the bilinear objective J_{xu}.

Bilinear objective J_{xu}: The bilinear objective function (from (2b)) is

$$J_{xu} = -ux = -(A_\mathrm{d}\sin(\omega_\mathrm{d} t) + \tilde{u}) \cdot (M\sin(\omega_\mathrm{d} t + \Psi) + M_0\tilde{u} - r). \tag{30}$$

Expanding and applying trigonometric identities gives:

$$\begin{aligned} J_{xu} = &-\frac{1}{2}A_\mathrm{d} M\cos(2\omega_\mathrm{d} t + \Psi) \\ &+ A_\mathrm{d} r\sin(\omega_\mathrm{d} t) - A_\mathrm{d} M_0\tilde{u}\sin(\omega_\mathrm{d} t) - M\tilde{u}\sin(\omega_\mathrm{d} t + \Psi) \\ &- \frac{1}{2}A_\mathrm{d} M\cos(\Psi) + \tilde{u}r - M_0\tilde{u}^2. \end{aligned} \tag{31}$$

As before, only the components at frequency ω_d (second line) are retained. Incorporating the high-pass filter $\mathcal{H}(\omega_{\mathcal{H}} < \omega_\mathrm{d})$ phase θ at ω_d and demodulating with $d_2(t,\psi)$, results in only the static terms:

$$\begin{aligned} \dot{\tilde{u}} = &-\frac{1}{2}A_\mathrm{d} M_0\tilde{u}\left(\cos(\psi - \theta) - \mathcal{O}(\geq 2\omega_\mathrm{d})\right) \\ &+ \frac{1}{2}A_\mathrm{d} r\left(\cos(\psi - \theta) - \mathcal{O}(\geq 2\omega_\mathrm{d})\right) \\ &- \frac{1}{2}M\tilde{u}\left(\cos(\psi - (\Psi + \theta)) - \mathcal{O}(\geq 2\omega_\mathrm{d})\right), \end{aligned} \tag{32}$$

where higher harmonic components at frequencies equal to or exceeding $2\omega_\mathrm{d}$, denoted by $\mathcal{O}(\geq 2\omega_\mathrm{d})$, are disregarded. Equating the obtained expression to zero ($\dot{\tilde{u}} = 0$) and rearranging gives

$$\tilde{u} = \frac{A_\mathrm{d} r\cos(\psi - \theta)}{A_\mathrm{d} M_0\cos(\psi - \theta) + M\cos(\psi - (\Psi + \theta))},$$

which simplifies to

$$\boxed{u^*_{xu} = \frac{r\cos(\psi - \theta)}{M_0\cos(\psi - \theta) + \left(M/A_\mathrm{d}\right)\cos(\psi - (\Psi + \theta))}}, \tag{33}$$

where $M = A_\mathrm{d}|G(\mathrm{j}\omega_\mathrm{d})|$ and $\Psi = \angle G(\mathrm{j}\omega_\mathrm{d})$ denote the magnitude and phase of the system frequency response evaluated at the dither excitation frequency ω_d. While both objectives J_{xx} and J_{xu} have equal steady-state optima in Section 4.1, it is remarkable that, in contrast to the result in (29), system dynamics in combination with the high-pass filter and demodulation phase influence not only the convergence dynamics, but also the final convergence result. A discussion of the results follows.

Discussion: Comparing (33) with (29) highlights a key distinction: for bilinear objectives, convergence properties depend on system dynamic properties at the dither excitation frequency, captured by the system gain $|G(\mathrm{j}\omega_\mathrm{d})|$ and phase Ψ, as well as by ESC-related phase shifts introduced through the high-pass filter phase θ and the demodulation phase ψ. In contrast, the state-dependent objective is insensitive to system dynamics. Achieving consistent convergence in the bilinear case therefore requires compensating both the system-induced magnitude and phase effects at the dither frequency, for instance, by selecting the demodulation phase such that $\psi = \theta$ and compensating the system gain accordingly.

This insight explains the common guideline of selecting the dither frequency in the slow-to-medium time scale, outside the system dynamic range (time-scale separation) [5]. In this regime, $|G(\mathrm{j}\omega_\mathrm{d})| \to M_0$ and $\Psi \to 0$, such that the influence of system dynamics effectively vanishes and (33) reduces to (29). Phase compensation using a phase-locked loop (PLL) has been proposed to correct the system phase [24,25]. However, the analysis presented in this section shows that compensating for the system gain at ω_d is also essential to ensure consistent convergence when operating beyond the slow time scale.

The presented analysis considers the open-loop mapping from the system input to the ESC gradient estimate, excluding the integrator dynamics. In the closed-loop ESC scheme, this gradient estimate, through the integrator, drives a slow adaptation of the quasi-steady input. Under the assumption of a sufficiently small convergence gain, the closed-loop dynamics can be interpreted via averaging arguments, whereby the integrator tracks the averaged gradient behavior. Consequently, local convergence of the closed-loop ESC system occurs toward the steady-state input at which the averaged gradient estimate vanishes, consistent with the expressions derived in the open-loop analysis.

The next section introduces approaches that remove the restrictive time-scale separation requirement, enabling reliable convergence at higher dither frequencies for bilinear objectives, thereby improving both consistency and convergence speed.

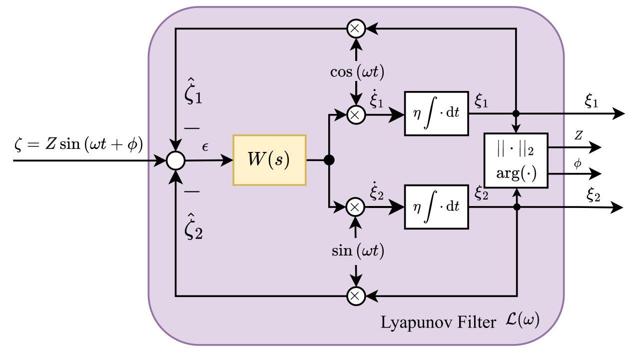

Fig. 7. Block diagram of the closed-loop Lyapunov amplitude demodulation filter (adapted from [30]). The estimation error ϵ is filtered by the transfer function $W(s)$, followed by mixing (multiplication) with two orthogonal sinusoids and scaling by the gain η. Integration yields the estimates of the quadrature, and in-phase states $\{\xi_1, \xi_2\}$, which are then fed back. The state estimates are then used to calculate the amplitude $Z(t)$ and phase $\phi(t)$ of the original periodic input signal $\zeta(t)$.

6. Solution paths

Previous analyses provided valuable insights into the convergence properties of state-dependent and bilinear objectives. Notably, unlike state-dependent objectives, the convergence of bilinear objectives requires strict adherence to the time-scale separation principle. The optimizing input for state-dependent objectives is given by u^*_{xx} in (29), while the bilinear objective converges to u^*_{xu} in (33) which – due to system dynamics and time-scale separation constraints – does not necessarily lead to the objective's steady-state optimum.

This section first provides preliminaries and then presents solution methods that achieve favorable convergence properties for bilinear objectives, inspired by the behavior of state-dependent objectives. The first section presents a signal-based phase- and gain-compensation method, while the second section proposes a numerical derivative approach to reconstruct the state-dependent objective.

6.1. Preliminaries – Lyapunov filter

This section outlines the principles of amplitude demodulation, followed by an explanation of the Lyapunov filter technique using synchronous detection [30]. Demodulation techniques enable the estimation of a signal's amplitude and phase for one or more frequency components. A generic amplitude-modulated signal is given by:

$$\zeta(t) = Z(t)\sin(\omega t + \phi(t)),$$

where the amplitude $Z(t)$ and phase $\phi(t)$ are to be estimated, and whose carrier frequency ω is known. Expanding the signal results in

$$\zeta(t) = \underbrace{Z(t)\cos(\omega t)\sin(\phi(t))}_{\text{quadrature component}} + \underbrace{Z(t)\sin(\omega t)\cos(\phi(t))}_{\text{in-phase component}},$$

which can be compactly written as

$$\zeta(t) = \varsigma(t)^\top\xi(t) = [\cos(\omega t)\ \sin(\omega t)]\begin{bmatrix} Z(t)\sin(\phi(t)) \\ Z(t)\cos(\phi(t)) \end{bmatrix}.$$

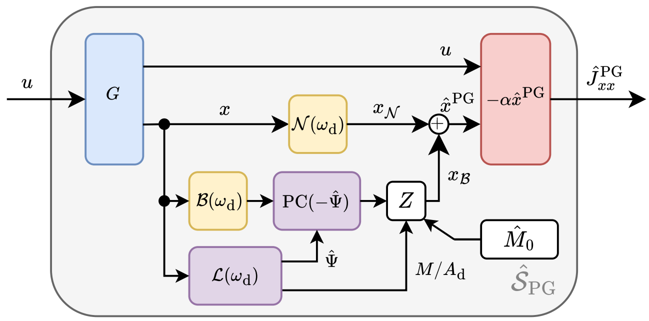

Fig. 8. Construction of the objective $\hat{J}_{xx}^{PG}$ to approximate the desired state-dependent objective function J_{xx}. The input signal u remains unaffected, while the system state x is decomposed in parallel using a notch and a band-pass filter. The band-pass filtered signal is phase-compensated by $\hat{\Psi}$, estimated by a Lyapunov filter. The signal is then amplitude-scaled and recombined with the notched signal to form $\hat{x}^{PG}$. Finally, the signal is multiplied by the input signal, resulting in a new phase-gain-compensated objective.

The aim of a demodulation strategy is to estimate the state vector $\xi(t) = \left[\xi_1, \xi_2\right]^\top$, where the state elements are the quadrature and in-phase components. From these, the amplitude and phase of the signal can be estimated as:

$$Z(t) = \sqrt{\xi_1^2(t) + \xi_2^2(t)}, \tag{34a}$$

$$\phi(t) = \arg(\xi_2(t) + \mathrm{j}\xi_1(t)). \tag{34b}$$

The Lyapunov filter [30], illustrated in Fig. 7, is a closed-loop estimator that uses a Lyapunov-based approach to estimate the parameters of linear parametric systems. It estimates the state vector $x(t)$ with feedback correction. Its state-space representation is:

$$\dot{\hat{\xi}}(t) = \eta\varsigma(t)^\top W(s)(\zeta(t) - \hat{\zeta}(t)), \tag{35a}$$

$$\hat{\zeta}(t) = \varsigma(t)\hat{\xi}(t), \tag{35b}$$

where $\hat{\xi}(t)$ is the estimated state vector, $W(s)$ is a strictly positive real (SPR) transfer ensuring closed-loop stability, and η is a gain parameter that calibrates the tracking bandwidth of the demodulator. The transfer function $W(s)$ is often chosen as unity for purely periodic signals, but may be chosen as a high-pass or a band-pass filter to remove the DC component in more general cases. Note that filtering the incoming signal affects the estimation bandwidth.

6.2. Phase-gain compensation

In earlier sections, steady-state optimizing inputs were derived for both state-dependent and bilinear objectives. This section introduces a method to reconstruct (29) from (33), ensuring convergence to the true steady-state optimum when using a bilinear objective. This is achieved through phase and gain compensation of the bilinear objective frequency content at the dither frequency, denoted by $\hat{(\cdot)}^{PG}$.

Examining (33) reveals that setting the demodulation phase ψ equal to the high-pass filter phase θ at the dither frequency results in

$$\hat{\hat{u}} = \frac{A_d r}{A_d M_0 + Z\left|G(\mathrm{j}\omega_d)\right|\cos(-\Psi)}, \tag{36}$$

where a new gain compensation factor $Z \in \mathbb{R}$ is introduced. To match (29), the last term in the denominator must equal M_0, leading to the definition:

$$Z = \frac{A_d \hat{M}_0}{\left|\widehat{G(\mathrm{j}\omega_d)}\right|\cos(-\hat{\Psi})}, \tag{37}$$

where $\hat{M}_0$ is the estimated system DC gain, and $\left|\widehat{G(\mathrm{j}\omega_d)}\right|$ and $\hat{\Psi}$ are the estimated gain and phase at the dither frequency, respectively. Note that $\left|\widehat{G(\mathrm{j}\omega_d)}\right| = \hat{M}/A_d$.

The three estimated elements described above are essential for the compensation method. The proposed method requires information about the system gain $\left|\widehat{G(\mathrm{j}\omega_d)}\right|$ and phase Ψ at the dither frequency. In this work, these estimates are obtained in a fully data-driven manner using a Lyapunov amplitude demodulation filter, which reconstructs the system response amplitude and phase at the dither frequency without requiring an explicit system model or parametric identification. As such, the proposed method remains consistent with the adaptive, model-free rationale of ESC. The only additional quantity assumed to be known or estimated independently is the DC system gain M_0, which can be obtained using simple steady-state experiments or standard system identification techniques and does not require detailed dynamic modeling.

From the bilinear formulation in Section 5.3, it is discovered that the state x only influences the last term in the denominator of (33). Therefore, the objective function is redefined as a phase-gain compensated approximation of the ideal objective:

$$\hat{J}_{xx}^{PG} = -\alpha(u)\hat{x}^{PG}, \tag{38}$$

where only $\hat{x}$ is modified from the original definition.

Fig. 8 illustrates the phase-gain compensation implementation. The state x is subject to parallel filters at the dither frequency: a notch filter $\mathcal{N}(\omega_d)$ removes the dither-periodic component, and a band-pass filter $\mathcal{B}(\omega_d)$ isolates the content at the dither frequency. In subsequent sections, filters centered around a given frequency are denoted using a single argument. That is, the filters defined above imply symmetric pass and rejection bands around ω_d with an appropriately chosen bandwidth. The band-pass filtered signal is phase compensated using a frequency-specific phase compensation method described in [31], with the phase estimate $\hat{\Psi}$ provided by the Lyapunov filter. This signal is amplitude-scaled by $\hat{M}_0/\left|\widehat{G(\mathrm{j}\omega_d)}\right|$ and recombined with the notched signal, leading to a reconstructed state $\hat{x}^{PG}$ that includes a phase-corrected and gain-scaled periodic component. These operations correspond to the definition of Z in (37). Note that, for an LTI system, the frequency response at a fixed dither frequency is time invariant for fixed model parameters.

Substituting the ESC input definition $u = A_d \sin(\omega_d t) + \tilde{u}$ and the reconstructed state $\hat{x}_{PG}$ into (38) yields:

$$\begin{aligned}\hat{J}_{xx}^{PG} &= -\alpha(u)\hat{x}^{PG}\\ &= -(A_d \sin(\omega_d t) + \tilde{u})\\ &\quad \cdot (\hat{M}_0 \sin(\omega_d t + (\Psi - \hat{\Psi})) + M_0\tilde{u} - r).\end{aligned} \tag{39}$$

Following the same quasi-static analysis as in Section 5.3, and assuming a consistent phase estimate $\hat{\Psi} = \Psi$ provided by the Lyapunov filter, the optimal input becomes:

$$\boxed{\hat{u}_{PG}^* = \frac{r}{M_0 + \hat{M}_0}}, \tag{40}$$

which closely resembles (29). The accuracy of convergence depends on the precision of the estimated quantities.

Fig. 9 shows ESC optimization results comparing the bilinear and phase-gain compensated objectives, J_{xu} and $\hat{J}_{xx}^{PG}$, respectively. The simulations are conducted at dither frequencies in distinct time scales to show the effectiveness of the proposed method. The definition of the dynamical system G and its parameters remains unchanged, with a steady-state gain of $M_0 = 1$. The red dashed line indicates the optimal steady-state input ($u^* = 3.5$). The ESC algorithm uses a convergence gain $\kappa = 0.02$ and a dither amplitude of $A_d = 4.0$ for all simulations, with a demodulation phase offset of $\psi = \theta$.

Both objectives converge around the true optimum at $\omega_d = 0.1$ rad/s (medium time scale). However, the bilinear objective shows inconsistent convergence at $\omega_d = 1.0$ rad/s (fast time scale), while the phase and gain-compensated objective maintains accurate convergence. As expected, higher dither frequencies yield faster convergence, demonstrating the method's ability to overcome time-scale separation limitations.

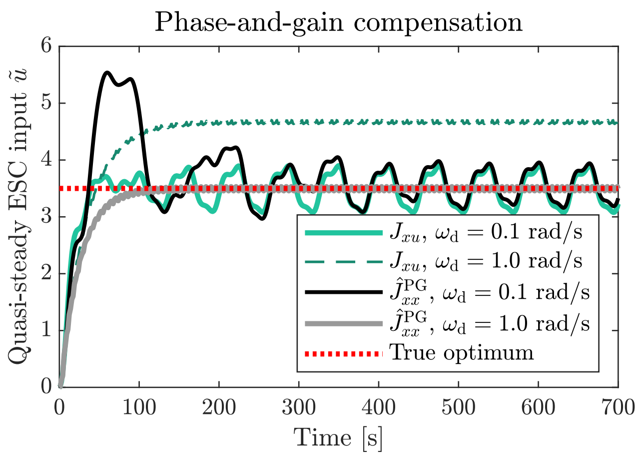

Fig. 9. ESC convergence using the bilinear objective J_{xu} and the phase-gain compensated objective $\hat{J}^{\mathrm{PG}}_{xx}$ at various dither frequencies. The dashed line indicates the true steady-state optimal input. At lower dither frequencies, both objectives converge near the optimum. At higher frequencies, the bilinear objective exhibits an offset with respect to the true optimum, whereas the PG-objective reduces the steady-state offset and improves convergence consistency around the true optimum.

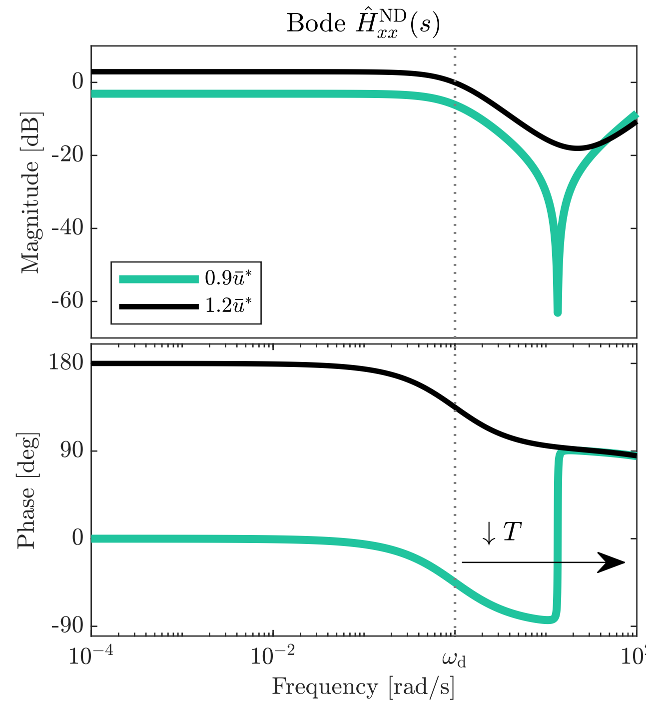

Fig. 10. Bode plot of the extended system with a state-derivative estimator based on a filtered numerical differentiator. Frequency responses are shown for two steady-state inputs, $\bar{u} = 0.9u^*$ and $\bar{u} = 1.2u^*$. The anti-resonance location depends on the filter time constant, as indicated in the phase plot. The system follows the trajectory of the state-dependent objective to the left of the anti-resonance and that of the bilinear objective to the right, enabling higher dither frequencies.

The phase-gain compensation method is intended to improve convergence consistency when operating at dither frequencies in the fast time scale, while the accuracy of the compensated convergence value depends on the quality of the estimated gain and phase information. The method relies on estimates of the system DC gain and the system gain and phase at the dither excitation frequency. In practice, inaccuracies in these estimated quantities and parametric uncertainties will primarily result in a bias of the steady-state convergence value with respect to the true steady-state optimum. Such inaccuracies will likely not lead to unstable behavior; instead, the ESC scheme continues to converge in a well-behaved manner to a steady operating point. In addition, noise on the estimated quantities increases the variance of the ESC gradient estimate, leading to larger residual oscillations around the converged result. A quantitative analysis of the impact of estimation errors is beyond the scope of this paper and will be considered in future work.

6.3. Objective reconstruction using numerical differentiation

This section provides a method for reconstructing the ideal state-dependent objective using filtered numerical differentiation. The method was originally proposed in [17] for ESC wind turbine controller optimization and is generalized in this work within the Wiener-type system framework.

To motivate the method, the state equation from (1) is reformulated as:

$$\gamma(x) = \alpha(u) - \tau\dot{x}, \tag{41}$$

with the input scaling coefficient taken as $c = 1$. Redefining the state equation provides the insight that reconstructing the state function $\gamma(x)$ requires knowledge of the system time constant τ and an estimate of the state derivative $\dot{x}$. Numerical derivative techniques obtain the latter component:

$$\hat{\gamma}^{\mathrm{ND}}(x) = \alpha(u) - \hat{\tau}\hat{\dot{x}}^{\mathrm{ND}}. \tag{42}$$

where the time constant is assumed to be known ($\hat{\tau} = \tau$), and $\hat{\dot{x}}^{\mathrm{ND}}$ is estimated via filtered numerical differentiation. The subscript $(\cdot)^{\mathrm{ND}}$ denotes numerical derivative-related variables.

The presented numerical-derivative-based reconstruction method assumes that the time constant τ of the dynamical system is known and accurately captured by $\hat{\tau}$. In practice, inaccuracies in the assumed time constant primarily result in a bias of the steady-state ESC convergence value with respect to the true steady-state optimum. It is also noted that a severe mismatch between the assumed time constant and the true system dynamics may adversely affect convergence behavior, including stability. Estimating or adapting $\hat{\tau}$ using simple dynamic experiments or system identification techniques is feasible and constitutes a relevant direction for future work, but is beyond the scope of the present paper.

The filtered numerical differentiator, based on [32], is defined by the following first-order system with unity gain:

$$\dot{\hat{x}}^{\mathrm{ND}} = f^{\mathrm{ND}}(\hat{x}^{\mathrm{ND}}, u^{\mathrm{ND}}) = \frac{1}{T}(u^{\mathrm{ND}} - \hat{x}^{\mathrm{ND}}), \tag{43a}$$

$$y = h^{\mathrm{ND}}(\hat{x}^{\mathrm{ND}}, u^{\mathrm{ND}}) = \dot{\hat{x}}^{\mathrm{ND}}, \tag{43b}$$

where T is the filter time constant, $\hat{x}^{\mathrm{ND}}$ is the filter internal state, and $u^{\mathrm{ND}} = x$ is the measured system state. The time constant determines the filter pole location and ensures the filter is proper. The filter time constant determines the trade-off between noise suppression and derivative accuracy: smaller values improve derivative accuracy but amplify high-frequency noise, while larger values reduce noise but degrade accuracy.

With the estimated state derivative, the ideal objective is approximated by substituting (42) and (43) into (2a):

$$\hat{J}^{\mathrm{ND}}_{xx} = -\hat{\gamma}^{\mathrm{ND}}(x)x = -(\alpha(u) - \tau\dot{\hat{x}}^{\mathrm{ND}})x. \tag{44}$$

To facilitate a similar analysis as done in the previous section, an augmented system $\hat{H}^{\mathrm{ND}}_{xx}$ is defined by the state vector $\left[x,\ \hat{x}^{\mathrm{ND}}\right]^\top$ and output $\hat{J}^{\mathrm{ND}}_{xx}$, with system matrices

$$\hat{A} = \begin{bmatrix} A & 0 \\ 1/T & -1/T \end{bmatrix}, \quad \hat{B} = \begin{bmatrix} B \\ 0 \end{bmatrix}$$
$$\hat{C} = \begin{bmatrix} \partial\hat{J}^{\mathrm{ND}}_{xx}/\partial x & \partial\hat{J}^{\mathrm{ND}}_{xx}/\partial x^{\mathrm{ND}} \end{bmatrix}, \quad \hat{D} = \begin{bmatrix} \partial\hat{J}^{\mathrm{ND}}_{xx}/\partial u \end{bmatrix},$$

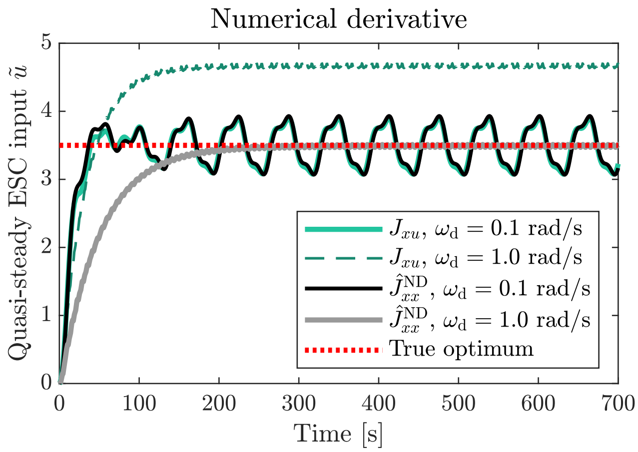

Fig. 11. Behavior of ESC with the bilinear objective J_{xu} and the numerical-derivative objective $\hat{J}^{ND}_{xx}$ for different dither frequencies. The dashed line marks the steady-state optimal input. At low dither frequency, both objectives converge near the optimum. At the high dither frequency, the bilinear objective deviates, while the ND-based objective remains closely aligned.

with the partial derivatives defined as

$$\frac{\partial \hat{J}^{ND}_{xx}}{\partial x} = \frac{\tau}{T}\left(2\bar{x} - \bar{\hat{x}}^{ND}\right) - \bar{\alpha},$$
$$\frac{\partial \hat{J}^{ND}_{xx}}{\partial \hat{x}^{ND}} = -\frac{\tau}{T}\bar{x}, \qquad \frac{\partial \hat{J}^{ND}_{xx}}{\partial u} = -\frac{\partial \alpha}{\partial u}\bar{x}.$$

Fig. 10 shows the Bode plot of $\hat{H}^{ND}_{xx}$ for operational conditions below and above the optimal input. The state and input function are chosen equal to the ones in Section 6.2, and the frequency responses are depicted for a filter time constant of $T = 10^{-3}$ s. When compared to the Bode plots in Figs. 5(a) and 5(b), the system now exhibits alternating phase behavior at $\omega_d = 1$ rad/s, which was previously absent. At low frequencies, the response aligns with the state-dependent objective, while at high frequencies (beyond the anti-resonance), it resembles the bilinear objective.

The filter approximates a pure differentiator as $T \to 0$, yielding an exact representation of J_{xx}. However, this increases sensitivity to noise. The choice of T thus balances favorable phase characteristics with robustness, enabling higher dither frequencies and partially relaxing time-scale separation constraints.

Fig. 11 illustrates the effectiveness of the method in a simulation. With $T = 10^{-3}$ s, ESC parameters $\kappa = 0.02$ and $A_d = 4.0$ are used. At a lower dither frequency, both objectives show similar performance and convergence near the true optimum. At high dither frequency, the bilinear objective deviates, while the ND-based objective maintains accurate convergence.

7. Conclusion

This paper investigated the convergence behavior of extremum seeking control (ESC) for dynamical systems with state-dependent and bilinear (state-input dependent) objectives. Although both formulations share the same steady-state optimum, the presented analyses demonstrate that their dynamic convergence properties can differ fundamentally. Frequency-domain analysis revealed that bilinear objectives may introduce alternating minimum-phase and nonminimum-phase behavior around the steady-state optimum, leading to sensitivity of the chosen dither frequency and potential convergence offsets unless strict time-scale separation is enforced. In contrast, state-dependent objectives retain favorable magnitude and phase characteristics and show consistent convergence even at higher dither frequencies. These observations were further supported by a signal-based analysis, which demonstrated that convergence under bilinear objectives depends explicitly on the system frequency response at the dither frequency and the phase shifts introduced by filtering and demodulation, whereas the state-dependent formulation yields convergence that is robust to such dynamical properties.

An important analytical outcome of this work is the explicit steady-state ESC convergence expression for bilinear objectives in the presence of system dynamics. This result shows that, unlike state-dependent objectives, the steady-state convergence point for bilinear formulations depends on the system frequency response at the dither excitation frequency, as well as on the dynamic characteristics introduced by the filtering and demodulation operations that are part of ESC. This analysis explains why bilinear objectives can exhibit inconsistent convergence behavior when the dither frequency is in the fast time scale, even though the steady-state optima are identical.

To address the limitations of bilinear objectives, two solution methods in the form of phase-gain compensation and numerical-derivative reconstruction are proposed. The phase-gain method can be implemented using input–output measurements and frequency-specific estimates, and reconstructs the desired convergence properties of state-dependent objectives. Second, a numerical-derivative-based method enables convergence at higher dither frequencies by approximating the ideal objective. It should be noted that this method requires access to the state or to a measured signal equivalent to the state used in the objective definition. Both methods were validated in simulation, showing improved convergence speed and more consistent behavior near the true optimum.

Overall, this work contributes to a deeper theoretical understanding of ESC objective formulation and provides tools for faster and more reliable optimization in systems where high-frequency dithering is beneficial or necessary. While the present analysis is restricted to SISO systems, extending the convergence analysis and the proposed compensation and reconstruction methods to multivariable extremum seeking control settings constitutes an interesting direction for future research.

CRediT authorship contribution statement

Sebastiaan Paul Mulders: Writing – review & editing, Writing – original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization. **Mario A. Rotea:** Writing – review & editing, Validation, Supervision. **Alexander Julian Gallo:** Writing – review & editing, Writing – original draft, Validation, Methodology, Investigation, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

[1] M. Hamza, Extremum control of continuous systems, IEEE Trans. Autom. Control 11 (2) (1966) 182–189, http://dx.doi.org/10.1109/TAC.1966.1098285.

[2] C.S. Draper, Y.T. Li, Principles of optimalizing control systems and an application to the internal combustion engine, ASME Publications, 1951.

[3] J. Creaby, Y. Li, J.E. Seem, Maximizing wind turbine energy capture using multivariable extremum seeking control, Wind Eng. 33 (4) (2009) 361–387, http://dx.doi.org/10.1260/030952409789685753.

[4] U. Ciri, M. Rotea, C. Santoni, S. Leonardi, Large-eddy simulations with extremum-seeking control for individual wind turbine power optimization, Wind. Energy 20 (9) (2017) 1617–1634, http://dx.doi.org/10.1002/we.2112.

[5] M. Krstic, H.-H. Wang, Stability of extremum seeking feedback for general nonlinear dynamic systems, Automatica 36 (4) (2000) 595–601, http://dx.doi.org/10.1016/S0005-1098(99)00183-1.

[6] K.B. Ariyur, M. Krstic, Real-Time Optimization By Extremum-Seeking Control, John Wiley & Sons, 2003, http://dx.doi.org/10.1002/0471669784.

[7] M. Rotea, Analysis of multivariable extremum seeking algorithms, in: Proceedings of the American Control Conference (ACC), Vol. 1, 2000, pp. 433–437, http://dx.doi.org/10.1109/ACC.2000.878937.

[8] N. Bizon, On tracking robustness in adaptive extremum seeking control of the fuel cell power plants, Appl. Energy 87 (10) (2010) 3115–3130, http://dx.doi.org/10.1016/j.apenergy.2010.04.007.

[9] D. Dochain, M. Perrier, M. Guay, Extremum seeking control and its application to process and reaction systems: A survey, Math. Comput. Simulation 82 (3) (2011) 369–380, http://dx.doi.org/10.1016/j.matcom.2010.10.022.

[10] W. Cai, A.N. Alhashim, I. N'Doye, T.M. Laleg-Kirati, An extremum seeking control based approach for alignment problem of mobile optical communication systems, IFAC-PapersOnLine 53 (2) (2020) 5368–5374, http://dx.doi.org/10.1016/j.ifacol.2020.12.1519.

[11] Y.-C. Chang, S. Gao, Extremum-seeking action selection for accelerating policy optimization, in: Proceedings of the IEEE International Conference on Robotics and Automation, ICRA, 2024, pp. 5141–5147, http://dx.doi.org/10.1109/ICRA57147.2024.10610197.

[12] Y. Tan, W.H. Moase, C. Manzie, D. Nešić, I.M. Mareels, Extremum seeking from 1922 to 2010, in: Proceedings of the 29th Chinese Control Conference, 2010, pp. 14–26.

[13] J. Baillieul, T. Samad, Encyclopedia of systems and control, Springer, 2021, http://dx.doi.org/10.1007/978-3-030-44184-5.

[14] Y. Tan, D. Nešić, I. Mareels, On the choice of dither in extremum seeking systems: A case study, Automatica 44 (5) (2008) 1446–1450, http://dx.doi.org/10.1016/j.automatica.2007.10.016.

[15] O. Trollberg, B. Carlsson, E.W. Jacobsen, Extremum seeking control of the CANON process - existence of multiple stationary solutions, J. Process Control 24 (2) (2014) 348–356, http://dx.doi.org/10.1016/j.jprocont.2013.11.007.

[16] M. Chioua, B. Srinivasan, M. Guay, M. Perrier, Dependence of the error in the optimal solution of perturbation-based extremum seeking methods on the excitation frequency, Can. J. Chem. Eng. 85 (4) (2007) 447–453, http://dx.doi.org/10.1002/cjce.5450850407.

[17] S.P. Mulders, A.J. Gallo, M.A. Rotea, Analysis of extremum seeking control for wind turbine torque controller optimization by aerodynamic and generator power objectives, in: Proceedings of the American Control Conference, ACC, 2024, pp. 1030–1037, http://dx.doi.org/10.23919/ACC60939.2024.10644899.

[18] M. Guay, D.J. Burns, A proportional integral extremum-seeking control approach for discrete-time nonlinear systems, in: Proceedings of the 54th IEEE Conference on Decision and Control, CDC, 2015, pp. 6948–6953, http://dx.doi.org/10.1109/CDC.2015.7403314.

[19] K.T. Atta, M. Guay, Fast proportional integral phasor extremum seeking control for a class of nonlinear system, IFAC-PapersOnLine 50 (1) (2017) 5724–5730, http://dx.doi.org/10.1016/j.ifacol.2017.08.1126.

[20] K.T. Atta, A. Johansson, T. Gustafsson, Extremum seeking control based on phasor estimation, Systems Control Lett. 85 (2015) 37–45, http://dx.doi.org/10.1016/j.sysconle.2015.08.010.

[21] K.T. Atta, M. Guay, Adaptive amplitude fast proportional integral phasor extremum seeking control for a class of nonlinear system, J. Process Control 83 (2019) 147–154, http://dx.doi.org/10.1016/j.jprocont.2018.10.006.

[22] T. van Keulen, R. van der Weijst, T. Oomen, Fast extremum seeking using multisine dither and online complex curve fitting, IFAC-PapersOnLine 53 (2) (2020) 5362–5367, http://dx.doi.org/10.1016/j.ifacol.2020.12.1228.

[23] W.H. Moase, C. Manzie, Fast extremum-seeking for Wiener–Hammerstein plants, Automatica 48 (10) (2012) 2433–2443, http://dx.doi.org/10.1016/j.automatica.2012.06.071.

[24] A.J. Peixoto, D. Pereira-Dias, A.F. Xaud, A.R. Secchi, Modelling and extremum seeking control of gas lifted oil wells, Proc. the 2nd IFAC Work. Autom. Control. Offshore Oil Gas Prod. 48 (6) (2015) 21–26, http://dx.doi.org/10.1016/j.ifacol.2015.08.004.

[25] A.J. Peixoto, D. Pereira-Dias, G.J. Tocaimasa, T.R. Oliveira, Fast extremum seeking for a class of non Hammerstein-wienner systems, in: Proceedings of the American Control Conference, ACC, 2017, pp. 5666–5671, http://dx.doi.org/10.23919/ACC.2017.7963837.

[26] M.A. Rotea, Logarithmic power feedback for extremum seeking control of wind turbines, IFAC-PapersOnLine 50 (1) (2017) 4504–4509, http://dx.doi.org/10.1016/j.ifacol.2017.08.381.

[27] J.-S. Deschênes, Demodulation considerations in extremum seeking control loops, in: Proceedings of the American Control Conference, ACC, 2012, pp. 3389–3395, http://dx.doi.org/10.1109/ACC.2012.6314824.

[28] M. Krstić, Performance improvement and limitations in extremum seeking control, Systems Control Lett. 39 (5) (2000) 313–326, http://dx.doi.org/10.1016/S0167-6911(99)00111-5.

[29] A.V. Oppenheim, Discrete-time Signal Processing, Pearson Education India, 1999.

[30] M.G. Ruppert, D.M. Harcombe, M.R. Ragazzon, S.R. Moheimani, A.J. Fleming, A review of demodulation techniques for amplitude-modulation atomic force microscopy, Beilstein J. Nanotechnol. 8 (2017) 1407–1426, http://dx.doi.org/10.3762/bjnano.8.142.

[31] A. Sharma, J.I.S. Hummel, S.P. Mulders, Single-blade individual pitch control with phase compensation for wind turbine periodic blade load reductions, in: Proceedings of the American Control Conference, ACC, 2025, pp. 2850–2855, http://dx.doi.org/10.23919/ACC63710.2025.11107742.

[32] IEEE, Ieee recommended practice for excitation system models for power system stability studies, 2016, http://dx.doi.org/10.1109/IEEESTD.2016.7553421, IEEE Std 421.5-2016.