Extremum Seeking Control with an Adaptive Gain Based On Gradient Estimation Error

Claus Danielson,* Scott A. Bortoff, and Ankush Chakrabarty†

Abstract

This paper presents an extremum-seeking control (ESC) algorithm with an adaptive step-size that adjusts the aggressiveness of the controller based on the quality of the gradient estimate. The adaptive step-size ensures that the integral-action produced by the gradient descent does not destabilize the closed-loop system. To quantify the quality of the gradient estimate, we present a batch least-squares (BLS) estimator with a novel weighting and show that it produces bounded estimation errors, where the uncertainty is due to the curvature of the unknown cost function. The adaptive step-size then maximizes the decrease of the combined plant and controller Lyapunov function for the worst-case estimation error. We prove that our ESC controller is input-to-state stable with respect to the dither signal. Finally, we demonstrate our ESC controller through five numerical examples; one illustrative, one practical, and three benchmarks.

1 Introduction

Extremum seeking control is a century-old [1] form of model-free adaptive control for the real-time optimization of dynamic systems. The objective of extremum-seeking control (ESC) is to drive the plant to an equilibrium that optimizes an unknown cost function. In this paper, we extend the proportional-integral extremum-seeking controller (PI-ESC) from [2–4] using an adaptive step-size that adjusts the aggressiveness of the controller based on the quality of the gradient estimate.

Most ESC controllers can be interpreted as gradient descent algorithms, wherein the controller follows a descent direction to the optimal. From the perspective of dynamic systems, gradient descent is integral-action. For instance, the standard gradient descent [2–7] wherein the updated set-point is the previous set-point minus a step-size times the gradient, has the dynamics of a discrete-time integrator. In classical ESC, the set-point is the continuous-time integral of the estimated gradients [8, 9]. Whether in discrete-time or continuous-time, it is a fundamental result from control-theory that integral-action can lead to instability. This issue is further complicated for non-linear systems. ESC controllers employ a variety of strategies to preserve stability. For instance, the inspiration for this paper [2–4] used a PI-ESC controller to preserved stability. In this paper, we introduce a integral gain matrix that ensured stability under the idealistic condition where the gradient is perfectly estimated. This integral gain represents the most aggressive step-size for the gradient descent that will not destabilize the plant. Our adaptive step-size attenuates this idealistic integral gain to preserve stability when the gradient estimate is imperfect.

---

*Department of Mechanical Engineering, University of New Mexico.

†Mitsubishi Electric Research Laboratories (MERL). Corresponding author. Email: achakrabarty@ieee.org. Phone: +1 (617) 758-6175.

Typically, strategies for promoting stability requires slowing the integral-action of the gradient descent, which can result in a slower convergence to the optimal. To accelerate convergence, other mechanisms are often added to the ESC algorithm. For example, dither adaptation has been explored in [2, 10], where the dither signal is made sufficiently small near the optimal solution so as to constrict the size of the uncertainty ball around the equilibrium state. Using the magnitude of the gradient for dither amplitude adaptation has been proposed in [11], and extended to dither adaptation using a super-twisting algorithm and higher-order sliding modes in [12]. The problem of removing excitation signals in a phasor-based ESC that incorporates PI control in the feedback path is discussed in [2]. The authors in [13] proposed a mechanism for reducing the dither amplitude in a traditional perturbation-based ESC strategy and provided a stability analysis. Methods for fast ESC convergence with high-frequency dither signals for systems with unknown dynamics is proposed in [14], and with unknown Hessians in [15]. Dither-free methods have also been explored in [16], and event-triggered mechanisms for fast convergence in [17]. In contrast, this paper focuses on accelerating convergence by using an adaptive step-size in the integral-action of the gradient descent. While dither perturbation has gained widespread attention, directly adapting the ESC controller gain is relatively uncommon, with a few exceptions, namely [18–20].

One of the main challenges of ESC is that the gradient of the unknown cost function must be estimated from data gathered while the system is in operation. Misestimating the gradient can exacerbate the stability issues introduced by the integral-action, especially when the gradient is over-estimated. This issue can be addressed by analyzing the ESC controller and gradient estimator in a common framework. For instance, ESC controllers often use an recursive least-squares (RLS) estimator to estimate the gradient of the cost function [2–5, 7]. This approach is attractive since the gradient estimator has state dynamics that can be analyzed in a common Lyapunov framework with the plant and ESC controller dynamics. The disadvantage of this approach is that it often results in conservatively slowing integral-action to allow the estimator to converge. In contrast, our ESC controller employs a batch least-squares (BLS) estimator with a novel weighting. The BLS estimator can be viewed as an operator that maps batches of collected data to gradient estimates. Thus, we do not need to consider its convergence rate in our analysis. Instead, we show that the novel weightings used in our BLS estimator produce bounded estimation errors. The adaptive step-size uses these bounds to maximize the decrease of the joint plant and controller Lyapunov function for the worst-case gradient estimation error. The potential advantage of this approach is that we can use more a aggressive integral-action on average without risking instability. This advantage is empirically demonstrated through our benchmark simulations. Another advantage of this approach is that it will facilitate future work based on more general gradient estimators, such as moving horizon estimators or set-based estimators.

For our ESC algorithm, persistently exciting data is necessary, but not sufficient, to accurately estimate the gradient of the cost. For accurate estimates, the data gathered must also be sufficiently close to equilibrium. To quantify the distance from equilibrium, we assume that our plant is instrumented to provide addition measurements beyond the cost, which is the only measurement used in most ESC problem formulations. Furthermore, we assume that the cost is a static function of these measurements. This is an admittedly strong assumption for ESC, but one that is consistent with many industrial systems which are heavily instrumented. Exploiting these additional sensor measurements to improve the convergence is a shrewd strategy. We show that accurate estimates of the gradient requires persistently exciting and sufficiently small perturbations of the input. Thus, the dither based acceleration methods described above can potentially be combined with our adaptive gain to further improve convergence.

The remainder of this paper is organized as follows. In Section 2, we define our ESC problem formulation. In Section 3, we present our ESC controller and prove its convergence. Finally, in Section 4 we demonstrate our ESC controller on five numerical example; one illustrative, one practical, and three benchmark.

Notation:

For a vector $v \in \mathbb{R}^n$ and square matrix $M \in \mathbb{R}^{n \times n}$, $\|v\|_M = \sqrt{v^\top M v}$ is the weighted 2-norm where the subscript is omitted for the identity matrix $\|v\| = \sqrt{v^\top v}$. For a square matrix $M \in \mathbb{R}^{n \times n}$, $\underline{\lambda}(M)$ and $\bar{\lambda}(M)$ denote its smallest and largest eigenvalues respectively and $\|M\| = \sup\{\|Mx\| : \|x\| \leq 1\}$ is the induced 2-norm. A function $\alpha : [0, \infty) \to [0, \infty)$ is class-$\mathcal{K}$ if $\alpha(0) = 0$ and it is strictly increasing. A function $\alpha : [0, \infty) \to [0, \infty)$ is class-$\mathcal{K}_\infty$ if it is class-$\mathcal{K}$ and $\lim_{a \to \infty} \alpha(a) = \infty$. A function $\beta : [0, \infty)^2 \to [0, \infty)$ is class-$\mathcal{KL}$ if $\beta(\cdot, t)$ is class-$\mathcal{K}$ $\forall t > 0$ and $\beta(r, \cdot)$ is continuous and strictly decreasing $\forall r > 0$. A system $x_{t+1} = f(x_t, d_t)$ is input-to-state stable if $\|x_t\| \leq \beta(\|x_0\|, t) + \gamma(\sup_t \|d_t\|)$ where $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$. A function f is $\mathcal{C}^n$ if the derivatives derivatives $f^{(1)}, \ldots, f^{(n)}$ exist and are continuous.

2 ESC Problem Statement

Consider the following discrete-time nonlinear system

$$x_{t+1} = f(x_t, u_t) \tag{1a}$$
$$y_t = g(x_t, u_t) \tag{1b}$$

where $x_t \in \mathbb{R}^{n_x}$ is the state, $u_t \in \mathbb{R}^{n_r}$ is the control input, and $y_t \in \mathbb{R}^{n_y}$ are measured output other than cost. We make the following assumptions about the plant (1).

Assumption 1 (Plant).

(a) The plant is controllable, observable, and Lipschitz continuous. Each input u corresponds a unique input-to-state stable (ISS) equilibrium state $\pi(u)$ where π is Lipschitz continuous.

(b) The output y_t tracks $y_t \to \bar{r}$ constant inputs $u_t = \bar{r}$.

Assumption 1 is admittedly a strong assumption, but one that is consistent with many industrial applications where ESC is applied to a closed-loop system with a well-designed controller and heavy instrumentation. When satisfied, this assumption can be used to improve the performance of ESC controllers. Assumption 1a means that the closed-loop system (1) is robustly stable. Thus, bounded perturbation of the input u_t cause bounded perturbation of the output y_t, allowing for safe exploration without risking instability. Assumption 1a is consistent with the assumptions made in other ESC literature e.g. [2–4].

Assumption 1b amounts to assuming that the steady-state map of the system is identity $g(\pi(u), u) = I$. This assumption is made for notational simplicity. The steady-state cost $\ell(u) = \mathcal{J}(g(\pi(u), u))$ with respect to the input u depends on the steady-state map $g(\pi(\cdot), \cdot)$. Without loss of generality, we can transform the inputs $u = \pi(r)$ to produce a plant (1) whose steady-state map $g(g^{-1}(\cdot)) = I$ is identity where g is invertible since both h and π are invertible. This simplifies the notation (but not the analysis) since $\nabla_r \mathcal{J}(g(\pi(r))) = \nabla_y \mathcal{J}$ instead of $\nabla_u \ell = \nabla_y \mathcal{J} \nabla_u h + \nabla_y \mathcal{J} \nabla_x h \nabla \pi$.

The **objective** of the ESC is to find a operating condition $u_t = \bar{r}^\star$ such that the plant (1) optimizes an unknown steady-state cost $\mathcal{J}(y)$. The optimal equilibrium is defined as

$$(\bar{y}^\star, \bar{u}^\star, \bar{x}^\star) = \arg\min \mathcal{J}(\bar{y}) \tag{2a}$$
$$\text{s.t. } \bar{x} = f(\bar{x}, \bar{u}) \tag{2b}$$
$$\bar{y} = g(\bar{x}, \bar{u}). \tag{2c}$$

We make the following assumptions about the cost $\mathcal{J} \in \mathcal{C}^2$.

Assumption 2 (Cost).

(a) The cost $\mathcal{J} \in \mathcal{C}^2$ has bounded curvature $\underline{H} \preceq \nabla^2 \mathcal{J} \preceq \bar{H}$.

(b) The cost $\mathcal{J} \in \mathcal{C}^2$ is bounded by class-$\mathcal{K}_\infty$ functions $\kappa_1(\|y - y^\star\|) \leq \mathcal{J}(y) \leq \kappa_2(\|y - y^\star\|)$ and its gradient $\nabla \mathcal{J}$ of the cost satisfies $\|\nabla \mathcal{J}(y)\| \geq \kappa_3(\|y - y^\star\|)$ for some class-$\mathcal{K}_\infty$ function κ_3.

Assumption 2a will be used to bound the estimation errors of our gradient estimator. This assumption holds if and only if the cost gradient $\nabla \mathcal{J} \in \mathcal{C}^1$ Lipschitz continuous i.e. $\underline{H} = -hI$ and $\bar{H} = hI$ implies $\|\nabla \mathcal{J}(y_1) - \nabla \mathcal{J}(y_2)\| \leq h\|y_1 - y_2\|$. However, our ESC algorithm can exploit more nuanced curvature bounds $\underline{H}, \bar{H}$, if available, to improve the convergence rate. Note that the bounds $\underline{H}$ and $\bar{H}$ are not required to be positive definite matrices. Thus, we are not assuming that the cost $\mathcal{J}$ is convex. Indeed, for two of our numerical examples, the cost will be non-convex.

Assumption 2b means that driving the cost gradient to zero $\nabla \mathcal{J} \to 0$ results in the output converging to the optimal $y_t \to y^\star$. This assumption will be used to prove the stability of the optimal equilibrium (2). If the cost $\mathcal{J}$ is convex (i.e. $0 \preceq \underline{H} \preceq \nabla^2 \mathcal{J}$) then Assumption 2a implies that Assumption 2b holds locally. However, this assumption can hold for non-convex costs $\mathcal{J}$, like those we will consider in our numerical examples.

3 Adaptive Gradient ESC Algorithm

Our extremum-seeking control (ESC) is given by

$$r_{t+1} = r_t + \begin{cases} -\alpha_t K \hat{\theta}_t & \text{if } \alpha_t \geq \underline{\alpha} \\ 0 & \text{otherwise} \end{cases} \tag{3a}$$

$$u_t = r_t + d_t \tag{3b}$$

where the state r_t of the controller is the current estimate of the optimal reference $\bar{r}^\star$ and the input u_t is the reference r_t plus a dither signal d_t. The step-size α_t and controller gain K will be described below. For $\alpha_t \geq \underline{\alpha}$, the ESC controller (3) is a discrete-time integral controller $r_{t+1} = r_t - \alpha_t K \hat{\theta}_t$.

The gradient $\theta_t = \nabla \mathcal{J}(r_t)$ of the cost function $\mathcal{J}(r_t)$ at the current reference set-point r_t is estimated by the following finite-horizon batch least-squares (BLS) estimator

$$\Lambda_t^{-1} = \frac{1}{N} \sum_{k=t-1}^{t-N} w_k \Delta y_k \Delta y_k^\top \tag{4a}$$

$$\hat{\theta}_t = \frac{\Lambda_t}{N} \sum_{k=t-1}^{t-N} w_k \Delta y_k \left(\Delta \mathcal{J}_k + \Delta y_k^\top \hat{H} \left(e_t - \tfrac{1}{2} \Delta y_k \right) \right) \tag{4b}$$

where $\Delta y_k = y_t - y_k$ and $\Delta \mathcal{J}_k = \mathcal{J}(y_t) - \mathcal{J}(y_k)$ are changes in the measurements of the output and cost, respectively, and the batch horizon $N \geq n_y$ is at least n_y. The existence of the inverse Λ_t of the information matrix (4a) requires that the output y_t of the plant (1) is persistently exciting, which is achieved using the dither d_t in the controller (3). The correction term $\Delta y_k^\top \hat{H}(e_t + \tfrac{1}{2}\Delta y_k)$ compensates for the tracking error $e_t = y_t - r_t$ and transients $\Delta y_k \neq 0$ where $\hat{H} = \tfrac{1}{2}(\bar{H} + \underline{H})$ is the *median* curvature $\nabla^2 \mathcal{J}$ of the cost $\mathcal{J}$. Without the measurements of the outputs y_t, the ESC controller (3) would need to be detuned to conservatively allow the plant (1) to settle near the equilibrium $\pi(r_t)$. If only a Lipschitz bound h on gradient $\nabla \mathcal{J}$ is known, then the correction term disappears $\Delta y_k^\top \hat{H}(e_t + \tfrac{1}{2}\Delta y_k) = 0$ since $\hat{H} = 0$ when $\underline{H} = -hI$ and $\bar{H} = hI$. For a convex cost with known Lipschitz bound h, the correction term is $\tfrac{h}{2}\Delta y_k^\top(e_t + \tfrac{1}{2}\Delta y_k)$. The **novel** weighting w_k is given by

$$w_k = \frac{1}{\tfrac{1}{2}\|\Delta y_k\|\|\Delta y_k\|_{\tilde{H}}\left(\|e_t\|_{\tilde{H}} + \tfrac{1}{2}\|\Delta y_k\|_{\tilde{H}}\right)} \tag{4c}$$

where $\tilde{H} = \bar{H} - \underline{H}$ bounds the range of curvature $\nabla^2 \mathcal{J}$ of the unknown cost $\mathcal{J}$. When the tracking errors $\|e_t\| \gg 1$ and output transients $\|\Delta y_k\| \gg 1$ are large, the weighting is small $w_k \ll 1$ indicating that the data-point $\{\Delta \mathcal{J}_k, \Delta y_k\}$ will not provide reliable information about the steady-state gradient $\nabla \mathcal{J}(r_t)$. When only a Lipschitz bound h on the gradient $\nabla \mathcal{J}$ is known, the weightings (4c) simplify

$$w_k = \frac{1}{h^2 \|\Delta y_k\|^2 \big(2\|e_t\| + \|\Delta y_k\|\big)}.$$

We will show that this weighting guarantees that the gradient estimation errors $\tilde{\theta}_t = \hat{\theta}_t - \nabla \mathcal{J}(r_t)$ are bounded.

Our **main contribution** is the adaptive step-size

$$\alpha_t = \max \left\{ 0, 1 - \frac{\|\Lambda_t^{\frac{1}{2}} K \hat{\theta}_t\|}{\|\hat{\theta}_t\|_K^2} \right\} \tag{5}$$

which dictates both the mode and the aggressiveness of the controller (3) based on the quality of the gradient estimate $\hat{\theta}_t$. If the step-size (5) is small $\alpha_t < \underline{\alpha}$, then the controller (3) enters the so called *exploration mode* where the state r_t of the controller (3) remains constant while the dither signal d_t probes the plant (1) to improve the gradient estimate. If $\alpha_t \geq \underline{\alpha}$ then the controller (3) enters the so called *exploitation mode* where it descends the estimated gradient $\hat{\theta}_t$ with $K = K^\top \succ 0$. Furthermore, the aggressiveness of this descent is dictated by the step-size (5). The max operator ensures that the step-size is non-negative and well-defined. If $\|\hat{\theta}_t\|_K^2 = 0$, then either we have perfectly misestimated the gradient $\tilde{\theta} = \nabla \mathcal{J}$ or perfectly estimated a zero gradient $\hat{\theta} = \nabla \mathcal{J} = 0$ (i.e. we are at optimal). In either case, the step-size (5) is zero since the controller should not step.

To better understand the intuition behind the adaptive step-size (5), consider the case where the controller gain K and estimator covariance Λ_t are balanced i.e. $K \approx kI$ and $\Lambda \approx \sigma^2 I$. Then, we can approximate the adaptive step-size (5) as

$$\alpha_t = \max \left\{ 0, 1 - \frac{\|\Lambda_t^{\frac{1}{2}}\| \|K \hat{\theta}_t\|}{\|\hat{\theta}\| \|K \hat{\theta}_t\|} \right\} \approx 1 - \frac{\sigma}{\mu}.$$

where σ/μ is the *noise-to-signal ratio* of the gradient estimate and $\mu = \|\hat{\theta}\|_K^2$ is the size of the descent direction. If the *noise-to-signal ratio* is small $\sigma/\mu \ll 1$, then $\alpha_t \approx 1$, allowing the ESC controller (3) to aggressively exploit the high-quality gradient estimate $\hat{\theta}_t$. Conversely, if the *noise-to-signal ratio* is large $\sigma/\mu \approx 1$ then the reduced step-size $\alpha_t \ll 1$ slows the gradient descent. Thus, the adaptive step-size (5) provides a reactive separation of time-scales between the controller (3) and estimator (4).

The positive definite controller gain $K = K^\top \succ 0$ of the ESC controller (3) must satisfy the matrix inequality

$$K - K\big(\bar{H} + \gamma I\big)K \succeq 0 \tag{6}$$

for some scalar γ. In Corollary 1 we will describe how to tuning of the controller gain (6) for a linear plant (1). If the plant (1) has trivial dynamics, then the gain (6) is $K = \frac{1}{2}\bar{H}$, which is the ideal choice for the (non-dynamic) optimization problem (2). For a dynamic plant, the gain (6) incorporates information about both the plant (1) and optimization problem (2) to improve convergence and prevent instabilty.

The following theorem proves that the ESC controller (3) converges to the reference $\bar{r}^\star$ that drives the plant (1) to a neighborhood of the optimal equilibrium (2).

Theorem 1. *Let Assumptions 1 and 2 hold. Let the dither d_t be persistently exciting and bounded $\|d_t\| \leq \delta$. Let K satisfy (6). Then the optimal equilibrium (2) is input-to-state stable for the closed-loop system (1) and (3)-(5).*

Theorem 1 says that the ESC controller (3) drives the plant (1) to a neighborhood of the optimal equilibrium (2) where the size of this neighborhood depends on the amplitude δ of the dither d_t. In practice, a vanishing dither [2, 10] can be used to provide convergence to the optimal, rather than only a neighborhood.

3.1 Proof of Theorem 1

In this section, we prove Theorem 1. First, we analyze the ESC controller (3) under the idealistic condition where the BLS estimator (4) is perfect $\hat{\theta}_t = \nabla \mathcal{J}(r_t)$ and thus the step-size (5) is maximal $\alpha_t = 1$. We will then examine how the adaptive step-size α_t can be used to make the ESC controller (3) robust to imperfect gradient estimates $\hat{\theta}_t \neq \nabla \mathcal{J}(r_t)$. Finally, we will show that our BLS estimator (4) satisfies our conditions for stability.

Proposition 1. *Let Assumptions 1 and 2 hold. Let K satisfy (6). Then the optimal equilibrium (2) is input-to-state stable (ISS) for the closed-loop system (1) and (3) where $\hat{\theta}_t = \nabla \mathcal{J}(r_t)$ and $\alpha_t = 1$.*

Proof. Define $\tilde{r}_t = r_t - \bar{r}^\star$ and $\tilde{x}_t = x_t - \pi(r_t)$. We will prove input-to-state stability using a candidate Lyapunov function of the form

$$V(\tilde{x}, \tilde{r}) = V_x(\tilde{x}) + V_r(\tilde{r}) \tag{7}$$

where V_x and V_r are candidate Lyapunov functions for the plant (1) and controller (3), respectively.

Since each constant equilibrium $\bar{x}^+ = \bar{x} = \pi(\bar{r})$ of the plant (1) is ISS by Assumption 1, there exists an ISS Lyapunov function V_x that satisfies

$$\underline{p}(\|\tilde{x}\|) \leq V_x(\tilde{x}) \leq \quad \bar{p}(\|\tilde{x}\|) \tag{8a}$$

$$V_x(f(x, \bar{u}) - \bar{x}^+) - V_x(\tilde{x}) \leq -q(\|\tilde{x}\|) + \sigma(\|d_t\|) \tag{8b}$$

according to the converse Lyapunov function theorem [21] where $\underline{p}$, $\bar{p}$, $q \in \mathcal{K}_\infty$ and $\sigma \in \mathcal{K}$. Here, the input $u_t = \bar{r} + d_t$ is dithered about the set-point $\bar{r}$. When the target equilibrium is varying $\bar{x}^+ \neq \bar{x}$, then the Lyapunov function (8) satisfies

$$\Delta V_x(\tilde{x}) = V_x(\tilde{x}^+) - V_x(x^+ - \bar{x}) + \underbrace{V_x(x^+ - \bar{x}) - V_x(\tilde{x})}_{\leq -q(\|\tilde{x}\|)}.$$

where the first-term is the increase of the Lyapunov function due to the changing set-point and the second-term decrease due to stability. We can assume without loss of generality that V_x is smooth [22], so that the first term above can be bounded by a quadratic

$$\Delta V_x(\tilde{x}) \leq \nabla V(\tilde{x})^\top \Delta \bar{x} + \tfrac{\rho}{2} \Delta \bar{x}^\top \Delta \bar{x}$$

where $\Delta \bar{x} = \bar{x}^+ - \bar{x}$ is the change in the equilibrium state and $\rho I \succeq \nabla^2 V_x$ is an upper-bound on the curvature of V_x. By Young's inequality $\nabla V^\top \Delta \bar{x} \leq \frac{1}{2\gamma_0} \nabla V_x^\top \nabla V_x + \frac{\gamma_0}{2} \Delta \bar{x}^\top \Delta \bar{x}$, we have

$$\Delta V_x(\tilde{x}) \leq -q + \frac{1}{2\gamma_0} \nabla V_x^\top P^{-1} \nabla V_x + \frac{\rho + \gamma_0}{2} \Delta \bar{x}^\top P \Delta \bar{x} \tag{9}$$

where $q_x^0 = q - \frac{1}{2\gamma_0}\nabla V_x^\top \nabla V_x \succ 0$ for an appropriate scaling of $\gamma_0 > 0$. Since π is Lipschitz continuous $\|\Delta \bar{x}\| \leq \ell_\pi \|\Delta r\|$, we obtain

$$\Delta V_x(\tilde{x}) \leq -q_x^0(\|\tilde{x}\|) + \tfrac{1}{2}\gamma \hat{\theta}^\top K^2 \hat{\theta}$$

where $\gamma = (\rho + \gamma_0)\ell_\pi^2$ and $\Delta r = K\hat{\theta}$.

A natural choice for the controller (3) Lyapunov function V_r is the cost function

$$V_r(\tilde{r}) = \mathcal{J}(\tilde{r} + \bar{r}^\star) - \mathcal{J}(\bar{r}^\star) \tag{10}$$

where $V_r(0) = 0$ by construction. By Assumption 2, the cost $\mathcal{J}$ is bounded above and below by class-$\mathcal{K}_\infty$ functions. By Taylor's theorem[1] and the controller dynamics (3), we have

$$\Delta V_r = V_r(\tilde{r}^+) - V_r(\tilde{r}) \leq \nabla \mathcal{J}(r)^\top \Delta r + \tfrac{1}{2}\Delta r^\top \bar{H} \Delta r$$
$$= -\hat{\theta}^\top K\hat{\theta} + \tfrac{1}{2}\hat{\theta}^\top K\bar{H}K\hat{\theta}$$

where $\hat{\theta} = \nabla \mathcal{J}$ and $\Delta r = K\hat{\theta}$ for $\alpha_t = 1$. Thus, the combined Lyapunov function (7) satisfies

$$\Delta V \leq -q_x^0(\|\tilde{x}\|) - \hat{\theta}^\top K\hat{\theta} + \tfrac{1}{2}\hat{\theta}^\top K(\gamma I + \bar{H})K\hat{\theta}$$
$$\leq -q_x^0(\|\tilde{x}\|) - \tfrac{1}{2}\hat{\theta}^\top K\hat{\theta}$$

where $K - \frac{1}{2}K(\gamma I + \bar{H})K \succeq \frac{1}{2}K$ by (6). By Assumption 2, the state $\tilde{r}$ of the controller (3) is bounded by the gradient $\kappa_3(\|\tilde{r}\|) \leq \|\theta\|$. Therefore, the combined Lyapunov function (7) is bounded by class-$\mathcal{K}_\infty$ functions and satisfies

$$\Delta V(\tilde{x}, \tilde{r}) \leq -q_x^0(\|\tilde{x}\|) - q_r^0(\|\tilde{r}\|) + \sigma(\|d_t\|)$$

where $q_x^0 = q - \frac{1}{2\gamma_0}\nabla V_x^\top \nabla V_x$ and $q_r^0(\tilde{r}) = \underline{\lambda}(K)\kappa_3(\|\tilde{r}\|)$ are class-$\mathcal{K}_\infty$ functions and $\sigma \in \mathcal{K}$. Thus, by Proposition 2.3 in [23] the optimal equilibrium (2) is ISS. $\square$

The Lyapunov function (7) defined in the proof of Proposition 1 will be used to prove Theorem 1. The proof of Proposition 1 uses similar arguments to the proof of Theorem 4.1 from [3]. However, our proof highlights the issue that without an appropriate controller gain (6), the integral-action of the gradient descent can destabilize the plant (1), even when the gradient is perfectly estimated $\hat{\theta} = \nabla \mathcal{J}$. As a quick aside, the following corollary shows how the derivation of the controller gain (6) for linear plants.

Corollary 1. *For a linear plant* (1), *Proposition 1 holds if the* ESC *controller gain K satisfies*

$$K - K(\bar{H} + B^\top (I-A)^{-\top} P Q^{-1} P (I-A)^\top B)K \succeq 0. \tag{11}$$

Proof. For a linear plant (1), we can use a quadratic $V_x(\tilde{x}) = \frac{1}{2}\tilde{x}^\top P\tilde{x}$ plant Lyapunov function (8) where P satisfies the Lyapunov equation $A^\top PA - P = -Q$ for some $Q \succ 0$. We can then use a matrix Γ, instead of a scalar γ_0, in Young's inequality (9) to obtain

$$\Delta V_x(\tilde{x}) = -\tilde{x}^\top Q\tilde{x} + \tfrac{1}{2}\tilde{x}^\top P\Gamma P\tilde{x} + \tfrac{1}{2}\Delta \bar{x}^\top (P + \Gamma^{-1})\Delta \bar{x}^\top$$

for some $\Gamma = \Gamma^\top \succ 0$ where $\nabla V_x = P\tilde{x}$ and $\nabla^2 V_x = P$. Following the argument of the proof of Proposition 1, we require $-Q + P\Gamma P \preceq 0$ for stability. Or equivalently $\Gamma \preceq P^{-1}QP^{-1}$. Thus,

$$\Delta V_x(\tilde{x}) \preceq \tfrac{1}{2}\Delta \bar{x}^\top (P + PQ^{-1}P)\Delta \bar{x}.$$

Finally, note that for a linear plant (1) the change in equilibrium state $\bar{x}$ satisfies $\Delta \bar{x} = (I-A)^{-1}B\Delta r = (I-A)^{-1}BK\hat{\theta}$. The remainder of the stability proof is identical to the proof of Proposition 1. $\square$

---

[1] Taylor's theorem, not a Taylor approximation.

The linear tuning in Corollary 1 can provide some insight for tuning the controller (3) gain (6) for a nonlinear plants (1), which is often challenging.

Next, we examine the robustness of the ESC controller (3) to imperfect gradient estimates $\hat{\theta}_t \neq \nabla \mathcal{J}(r_t)$. We will consider gradient estimation errors $\tilde{\theta}_t := \hat{\theta}_t - \nabla \mathcal{J}(r_t)$ that are slightly smaller $\|\tilde{\theta}\|_K \leq (1-\underline{\alpha})\|\hat{\theta}\|_K$ than the estimated gradient $\hat{\theta}_t$

$$\tilde{\Theta}_t^{\max} = \left\{ \tilde{\theta} \in \mathbb{R}^{n_y} : \|\tilde{\theta}\|_K \leq (1-\underline{\alpha})\|\hat{\theta}_t\|_K \right\} \tag{12}$$

where $0 < \underline{\alpha} \ll 1$. The following corollary shows that the optimal equilibrium (2) remains ISS for gradient estimation errors $\tilde{\theta}_t = \hat{\theta}_t - \nabla \mathcal{J}(r_t)$ that satisfy the bound $\tilde{\theta}_t \in \tilde{\Theta}_t \subseteq \tilde{\Theta}_t^{\max}$.

Corollary 2. *Let Assumptions 1 and 2 hold. Let the set $\tilde{\Theta}_t$ of gradient estimation errors $\tilde{\theta}_t = \hat{\theta}_t - \nabla \mathcal{J}_t \in \tilde{\Theta}_t$ satisfy the bound $\tilde{\Theta}_t \subseteq \tilde{\Theta}_t^{\max}$. Let K satisfy (6) and the step-size α_t satisfy*

$$\alpha_t^\star = \max \left\{ 0, 1 - \max_{\tilde{\theta} \in \tilde{\Theta}_t} \frac{\tilde{\theta}^\top K \hat{\theta}_t}{\|\hat{\theta}_t\|_K^2} \right\}. \tag{13}$$

Then the optimal equilibrium (2) is ISS for the closed-loop system (1) and (3)

$$\Delta V(\tilde{x}, \tilde{r}) \leq -q_x^1(\|\tilde{x}\|) - q_r^2(\|\tilde{r}\|) + \sigma(\|d_t\|) \tag{14}$$

where $q_x^1, q_r^2 \in \mathcal{K}_\infty$ and $\sigma \in \mathcal{K}$.

Proof. For imperfect gradient estimates $\hat{\theta} \neq \nabla \mathcal{J}$, the Lyapunov function (7) from Proposition 1 satisfies

$$\Delta V \leq -q_x^0 - \alpha_t \nabla \mathcal{J}^\top K \hat{\theta} + \tfrac{1}{2}\alpha_t^2 \hat{\theta}^\top K \hat{\theta} + \sigma(\|d_t\|) \tag{15}$$

where the step-size $\alpha_t \geq 0$ is a design variable we can choose to promote stability while the worst-case gradient estimation error $\tilde{\theta}_t \in \tilde{\Theta}_t$ will try to prevent stability. Since the worst-case estimation error $\tilde{\theta}_t \in \tilde{\Theta}_t$ will depend on our choice of step-size α_t, we have the following two-player zero-sum game

$$\min_{\alpha_t \geq 0} \max_{\tilde{\theta}_t \in \tilde{\Theta}_t} \; -q_x^0 - \alpha_t(\hat{\theta} - \tilde{\theta})^\top K \hat{\theta} + \tfrac{1}{2}\alpha_t^2 \hat{\theta}^\top K \hat{\theta} + \sigma(\|d_t\|) \tag{16}$$

where $\nabla \mathcal{J} = \hat{\theta} - \tilde{\theta}$. Here, the adversary has the advantageous position of selecting the worst-case gradient estimation error $\tilde{\theta}_t \in \tilde{\Theta}_t$ based on our choice of step-size $\alpha_t \geq 0$. However, for this particular game (16) the optimal strategy $\tilde{\theta}_t^\star(\alpha_t)$ for the gradient estimation error happens to be independent of the step-size since

$$\tilde{\theta}_t^\star = \arg \max_{\tilde{\theta}_t \in \tilde{\Theta}_t} \; \alpha_t \hat{\theta}_t^\top K \tilde{\theta}$$

where the non-negative step-size $\alpha_t \geq 0$ only scales the linear cost $\hat{\theta}^\top K$, but does not change its direction. The optimal strategy for the step-size $\alpha_t^\star$ is given by the following scalar quadratic program

$$\alpha_t^\star = \arg \min_{\alpha_t \geq 0} \; -q_x^0 - \alpha_t(\hat{\theta} - \tilde{\theta}^\star(\alpha_t))^\top K \hat{\theta} + \tfrac{1}{2}\alpha_t^2 \hat{\theta}^\top K \hat{\theta}$$

$$= \arg \min_{\alpha_t \geq 0} \; \alpha_t\big(\hat{\theta}^\top K \tilde{\theta}^\star - \|\hat{\theta}\|_K^2\big) + \tfrac{1}{2}\alpha_t^2 \|\hat{\theta}\|_K^2$$

where $\tilde{\theta}^\star(\alpha_t) = \max_{\tilde{\theta} \in \tilde{\Theta}_t} \tilde{\theta}^\top K \hat{\theta}_t$ is independent of α_t. The step-size (13) is the explicit optimal solution of this parametric quadratic program where $\tilde{\theta}_t^\star$ is the parameter. Thus, (13) is the game-theoretic optimal

step-size α_t for bounded gradient estimation errors $\tilde{\theta}_t \in \tilde{\Theta}_t$. The resulting change (15) of the Lyapunov function (7) satisfies

$$\Delta V = \Delta V_x + \Delta V_r \leq -q_x^0 - (\alpha_t^\star)^2 \|\hat{\theta}\|_K^2 + \sigma(\|d_t\|).$$

To prove ISS, we will show that $(\alpha_t^\star)^2 \|\hat{\theta}\|_K^2$ is bounded by a class-$\mathcal{K}_\infty$ function $q_r^1(\|\tilde{r}\|)$ of the controller (3) state $\tilde{r}$. Since $\tilde{\Theta}_t \subseteq \tilde{\Theta}_t^{\max}$, we have

$$\alpha_t^\star \geq 1 - \max_{\tilde{\theta} \in \tilde{\Theta}_t} \frac{\tilde{\theta}^\top K \hat{\theta}_t}{\|\hat{\theta}_t\|_K^2} \geq 1 - \max_{\tilde{\theta} \in \tilde{\Theta}_t} \frac{\|\tilde{\theta}_t\|_K \|\hat{\theta}_t\|_K}{\|\hat{\theta}_t\|_K^2} = \underline{\alpha}$$

where the first inequality follows from the step-size (13), the second inequality is the Cauchy-Schwarz inequality, and the last inequality follows from the definition (12) of the set $\tilde{\Theta}_t^{\max}$. Thus,

$$(\alpha_t^\star)^2 \|\hat{\theta}\|_K^2 \geq \underline{\alpha}^2 \|\hat{\theta}\|_K^2$$

where $\underline{\alpha} > 0$. Next, we need to bound the norm $\|\hat{\theta}\|_K^2$ of the estimated gradient $\hat{\theta}_t$ by the norm $\|\theta\|_K^2$ of the actual gradient $\theta_t = \nabla \mathcal{J}(r_t)$. Since $\theta_t = \hat{\theta}_t - \tilde{\theta}_t$, we have

$$\|\theta\|_K = \|\hat{\theta} - \tilde{\theta}\|_K \leq \|\hat{\theta}\|_K + \|\tilde{\theta}\|_K$$
$$\leq \|\hat{\theta}\|_K + (1 - \underline{\alpha})\|\hat{\theta}\|_K \leq 2\|\hat{\theta}\|_K$$

where the first inequality is the triangle inequality and the second inequality follows from the definition (12) of the set $\tilde{\Theta}_t^{\max}$. Rearranging terms and squaring yields $\underline{\alpha}^2 \|\hat{\theta}\|_K^2 \geq \frac{1}{4}\underline{\alpha}^2 \|\theta\|_K^2$. Finally, recall from the proof of Proposition 1 that

$$\|\theta\|_K^2 \geq \underline{\lambda}(K)\kappa_3(\|\tilde{r}\|),$$

where $\kappa_3(\|\tilde{r}\|) \in \mathcal{K}_\infty$. Therefore, we conclude that (14) holds where $q_x^1(\|\tilde{x}\|) = q_x^0(\|\tilde{x}\|) \in \mathcal{K}_\infty$ and

$$q_r^2(\|\tilde{r}\|) = \frac{1}{4}\underline{\alpha}^2 \underline{\lambda}(K)\kappa_3(\|\tilde{r}\|) \in \mathcal{K}_\infty.$$

$\square$

Corollary 2 shows that the adaptive step-size (13) makes our ESC controller (3) robust to bounded (12) gradient estimation errors $\tilde{\theta}_t = \hat{\theta}_t - \nabla \mathcal{J}_t \in \tilde{\Theta}_t$. The adaptive step-size (13) ensures that the Lyapunov function (7) decreases (14) for the worst-case gradient estimation errors $\tilde{\theta}_t \in \tilde{\Theta}_t$. When the bounds $\tilde{\Theta}_t$ on the gradient estimation errors $\tilde{\theta}_t \in \tilde{\Theta}_t$ are tight, the step-size (13) is large and the controller (3) is aggressive. Interestingly, the worst-case gradient estimation error $\tilde{\theta}_t^\star$ does not try to alter the descent-direction $K\hat{\theta}$, but rather amplifies the descent-direction $\hat{\theta} = \frac{1}{\underline{\alpha}}\theta$. This can cause the integral-action of the ESC controller (3) to overshoot the optimal, potentially leading to unstable oscillation.

Later in Lemma 1, we will derive the specific bounds $\tilde{\Theta}_t^{\mathrm{BLS}}$ on the gradient estimation errors produce by the BLS estimator (4). We will also connect the game-theoretic step-size (13) with the step-size (5) used in the ESC controller (3) in Corollary 4.

The set (12) is the largest set $\tilde{\Theta}_t^{\max}$ of gradient estimation errors $\tilde{\theta}$ for which the ESC controller (3) can decrease the Lyapunov function (7). Thus, the controller remains in the exploration mode i.e. $\alpha_t^\star \geq \underline{\alpha}$. When the bounds $\tilde{\Theta}_t$ on the estimation errors are too large $\tilde{\Theta}_t \not\subseteq \tilde{\Theta}_t^{\max}$, the ESC controller (3) enters the exploration mode i.e. $\alpha_t^\star = 0$. The following trivial corollary shows that optimal equilibrium (2) remains stable (but not ISS nor asymptotically stable) when the ESC controller (3) is in exploration mode $\alpha_t^\star = 0$.

Corollary 3. *Let Assumptions 1 and 2 hold. Let K satisfy (6) and the step-size satisfy $\alpha_t^\star = 0$. Then the plant (1) is* ISS *with respect to the dither signal d_t*

$$\Delta V_x(\tilde{x}) \leq -q_x^2(\|\tilde{x}\|) + \sigma(\|d_t\|), \tag{17a}$$

and the controller (3) is stable

$$\Delta V_r(\tilde{r}) = 0, \tag{17b}$$

where $q_x^2 \in \mathcal{K}_\infty$ and $\sigma \in \mathcal{K}$.

Proof. Since the controller (3) state $r_{t+1} = r_t$ is constant in the exploration mode, the controller Lyapunov function is constant (17b). The decrease condition (17a) for the plant Lyapunov function follows from the ISS assumption (8). □

Corollary 3 shows that the divergence from the optimal equilibrium (2) is bounded for bounded dithers d_t. This allows the dither d_t to safely probe the system (1) to gather exciting data and reduce the gradient estimation errors $\tilde{\Theta}_t$. However, it does not provide convergence towards the optimal equilibrium (2). For convergence, we need the ESC controller (3) to interminably enter the exploitation mode $\alpha_t^\star \geq \underline{\alpha}$. This requires that the bounds $\tilde{\Theta}_t^{\text{BLS}}$ on the gradient estimation errors produce by the BLS estimator (4) are sufficient small (12). The following lemma establishes bounds $\tilde{\Theta}_t^{\text{BLS}}$ on the gradient estimation errors produce by the BLS estimator (4).

Lemma 1. *Let Assumptions 1 and 2 hold. The estimation errors $\tilde{\theta}_t = \hat{\theta}_t - \nabla \mathcal{J}_t$ produced by the BLS estimator (4) are contained in the set*

$$\tilde{\Theta}_t^{\text{BLS}} = \left\{ \tilde{\theta} : \left\| \Lambda_t^{-1} \tilde{\theta} \right\|^2 \leq 1 \right\}. \tag{18}$$

Proof. To derive the set-bound (18) on the gradient estimation error $\tilde{\theta}_t$, we will express the gradient $\theta = \nabla \mathcal{J}$ as a linear regression $\Delta \mathcal{J}_k = \Delta y_k^\top \theta + \omega_k$ where the *noise* ω_k is set bounded, rather than stochastic.

According to Taylor's theorem[2], the cost $\mathcal{J}$ satisfies

$$\Delta \mathcal{J}(y_k) = \Delta y_k^\top \nabla \mathcal{J}(y_t) + \tfrac{1}{2} \Delta y_k^\top \nabla^2 \mathcal{J}(z_1) \Delta y_k \tag{19}$$

where the unknown curvature $H_1 = \nabla^2 \mathcal{J}(z_1)$ is evaluated at an unknown point $z_1 = \mu_1 y_t + (1 - \mu_1) y_k$ for some $\mu_1 \in [0, 1]$. The desired gradient $\nabla \mathcal{J}(r_t)$ evaluated at the current reference r_t is related to the gradient $\nabla \mathcal{J}(y_t)$ evaluated at the current output y_t by the mean-value theorem

$$\theta_t = \nabla \mathcal{J}(r_t) = \nabla \mathcal{J}(y_t) - \nabla^2 \mathcal{J}(z_2)^\top (r_t - y_t) \tag{20}$$

where $e_t = r_t - y_t$ is the tracking error and the unknown curvature $H_2 = \nabla^2 \mathcal{J}(z_2)$ is again evaluated at an unknown point $z_2 = \mu_2 y_t + (1 - \mu_2) r_t$ for some $\mu_2 \in [0, 1]$. Combining the non-approximations (19) and (20), yields

$$\Delta \mathcal{J}_k = \Delta y_k^\top \theta_t + \underbrace{\tfrac{1}{2} \Delta y_k^\top H_1 \Delta y_k}_{\omega_1} - \underbrace{e_t^\top H_2 \Delta y_k}_{\omega_2} \tag{21}$$

where $\theta_t = \nabla \mathcal{J}(r_t)$, $\Delta \mathcal{J}_k = \mathcal{J}(y_k) - \mathcal{J}(y_t)$ and $\Delta y_k = y_k - y_t$. The non-approximation (21) says that the cost $\mathcal{J}$ can be written as a linear regression $\Delta \mathcal{J}_k = \Delta y_k^\top \theta + \omega_k$ where the unknown error term $\omega = \omega_1 - \omega_2$

---

[2]Note that we use Taylor's theorem, which is exact, not a Taylor approximation, which is an approximation.

accounts for the nonlinearity. We will show that the error term ω is bounded since the curvature $\nabla^2 \mathcal{J}$ of the cost $\mathcal{J}$ is bounded $\underline{H} \preceq \nabla^2 \mathcal{J} \preceq \bar{H}$.

From the definition of positive definite matrices, the quadratic-form $\omega_1 = \frac{1}{2}\Delta y_k^\top H_1 \Delta y_k$ is contained in the line interval

$$\Omega_k^1 = \left\{ \omega_1 \in \mathbb{R} : \tfrac{1}{2}\Delta y_k^\top \underline{H} \Delta y_k \le \omega_1 \le \tfrac{1}{2}\Delta y_k^\top \bar{H} \Delta y_k \right\}$$

$$= \tfrac{1}{2}\|\Delta y_k\|_{\hat{H}}^2 + \tfrac{1}{4}\|\Delta y_k\|_{\tilde{H}}^2 \Big[-1, 1 \Big]$$

where $\hat{H} = \frac{1}{2}(\bar{H} + \underline{H})$ is the median curvature and $\tilde{H} = \bar{H} - \underline{H}$ is the range of curvature. Deriving the bounds Ω_k^2 on the nonlinearity $\omega_2 = e_t^\top H_2 \Delta y_k$ is more complicated since it is not a quadratic-form. Since the linear function $f(H) = e_t^\top H \Delta y_k$ is continuous, the image $\Omega_k^2 = f(\mathcal{H})$ of the connected set $\mathcal{H} = \{H : \underline{H} \preceq H \preceq \bar{H}\}$ is connected. Furthermore, since $\omega_2 \in \mathbb{R}$ is a scalar, this set Ω_k^2 is a line interval, specifically

$$\Omega_k^2 = \left\{ \omega_2 : \min_{\underline{H} \preceq H \preceq \bar{H}} \tfrac{1}{2}\mathrm{Tr}(CH) \le \omega_2 \le \max_{\underline{H} \preceq H \preceq \bar{H}} \tfrac{1}{2}\mathrm{Tr}(CH) \right\} \tag{22}$$

where

$$C = \Delta y_k e_t^\top + e_t \Delta y_k^\top$$

is the symmetric cost matrix. In other words, we can find the bounds on ω_2 by solving two semi-definite programs (22) for the lower and upper bounds on the line interval. According to Theorem 2.2 from [24], these semi-definite programs (22) have a closed-from solution, specifically

$$\min_{\underline{H} \preceq H \preceq \bar{H}} \tfrac{1}{2}\mathrm{Tr}(CH) = \tfrac{1}{2}\mathrm{Tr}^-\big(\tilde{H}^{\frac{1}{2}}C\tilde{H}^{\frac{1}{2}}\big) + \tfrac{1}{2}\mathrm{Tr}\big(C\underline{H}\big)$$

where Tr^- is the trace of the projection of a matrix into the negative semi-definite cone i.e. the sum of its negative eigenvalues. The rank-2 matrix $\tilde{H}^{\frac{1}{2}}C\tilde{H}^{\frac{1}{2}}$ has exactly one negative eigenvalue

$$\lambda_- = e_t^\top \tilde{H} \Delta y_k - \|\Delta y_k\|_{\tilde{H}}\|e_t\|_{\tilde{H}}$$

. Thus,

$$\min_{\underline{H} \preceq H \preceq \bar{H}} \tfrac{1}{2}\mathrm{Tr}(CH) = \tfrac{1}{2}\Delta y_k^\top \tilde{H} e_t - \tfrac{1}{2}\|\Delta y_k\|_{\tilde{H}}\|e_t\|_{\tilde{H}} + e_t^\top \underline{H}\Delta y_k$$

$$= \Delta y_k^\top \hat{H} e_t - \tfrac{1}{2}\|\Delta y_k\|_{\tilde{H}}\|e_t\|_{\tilde{H}}.$$

Similarly, we can obtain the upper-bound

$$\omega_2 \le \Delta y_k^\top \hat{H} e_t + \|\Delta y_k\|_{\tilde{H}}\|e_t\|_{\tilde{H}}.$$

Thus,

$$\Omega_k^2 = \Delta y_k^\top \hat{H} e_t + \tfrac{1}{2}\|\Delta y_k\|_{\tilde{H}}\|e_t\|_{\tilde{H}} \Big[-1, 1 \Big]$$

Therefore, the total error $\omega = \omega_1 - \omega_2$ is contained in the line interval

$$\Omega_k = -\Delta y_k^\top \hat{H}\big(e_t - \tfrac{1}{2}\Delta y_k\big) + \frac{1}{w_k\|\Delta y\|}\Big[-1, 1 \Big].$$

where the weighting w_k was defined in (4c). Thus, cost (21) can be rewritten

$$\Delta \mathcal{J}_k + \Delta y_k^\top \hat{H}\left(e_t - \tfrac{1}{2}\Delta y_k\right) = \Delta y_k^\top \theta_t + \frac{1}{w_k\|\Delta y_k\|}\nu_k \tag{23}$$

where $\nu_k \in [-1, 1]$ is the normalized *noise*. Substituting (23) into the BLS (4) yields

$$\hat{\theta}_t = \Lambda_t \underbrace{\left(\frac{1}{N}\sum_{k=t-1}^{t-N} w_k \Delta y_k \Delta y_k^\top\right)}_{\Lambda_t^{-1}}\theta_t + \Lambda_t \frac{1}{N}\sum_{k=t-1}^{t-N}\frac{\Delta y_k}{\|\Delta y_k\|}\nu_k.$$

Thus, by definition of the estimation error $\tilde{\theta}_t = \hat{\theta}_t - \theta_1$, the bound (18) holds since

$$\left\|\Lambda_t^{-1}\tilde{\theta}_t\right\| = \left\|\frac{1}{N}\sum_{k=t-1}^{t-N}\frac{\Delta y_k}{\|\Delta y_k\|}\nu_k\right\| \le \frac{1}{N}\sum_{k=t-1}^{t-N}|\nu_k| \le 1$$

where the first inequality is the triangle inequality and the second inequality follows from $|\nu_k| \le 1$ since $\nu_k \in [-1, 1] \subset \mathbb{R}$. $\qquad\square$

Lemma 1 describes a (possibly degenerate) ellipsoidal set $\tilde{\Theta}_t^{\mathrm{BLS}}$ that bounds the gradient estimation errors produce by the BLS estimator (4). The following lemma shows that for persistently exciting and local data $\{\mathcal{J}_k, y_k\}_{t=k}^t$, the ellipsoid (18) is non-degenerate and eventually satisfies the bounds $\tilde{\Theta}_t^{\mathrm{BLS}} \subseteq \tilde{\Theta}_t^{\max}$, allowing the ESC controller (3) to return to the exploitation mode.

Lemma 2. *Let Assumptions 1 and 2 hold. Let $\nabla\mathcal{J}(r_t) \neq 0$. Let the dither d_t be persistently exciting and bounded $\|d_t\| \le \delta$. Then there exists a finite time $T < \infty$ and non-zero dither amplitude $\delta > 0$ such that (18) satisfies $\tilde{\Theta}_{t+T}^{\mathrm{BLS}} \subseteq \tilde{\Theta}_{t+T}^{\max}$.*

Proof. For notational simplicity, we will drop the time indices.

First, we will transform the desired condition $\tilde{\Theta}^{\mathrm{BLS}} \subseteq \tilde{\Theta}^{\max}$ into a more readily verifiable form. From the definition (12) of $\tilde{\Theta}^{\max}$, the condition $\tilde{\Theta}^{\mathrm{BLS}} \subseteq \tilde{\Theta}^{\max}$ holds if and only if $\|\tilde{\theta}\|_K \le (1-\alpha)\|\hat{\theta}\|_K$ for all $\tilde{\theta} \in \tilde{\Theta}^{\mathrm{BLS}}$. Substituting $\hat{\theta} = \theta + \tilde{\theta}$ and expanding the norm $\|\hat{\theta}\|_K^2 = \|\theta + \tilde{\theta}\|_K^2$, we have the following equivalent condition

$$\alpha(2-\alpha)\|\tilde{\theta}\|_K^2 - 2(1-\alpha)^2\theta^\top K\tilde{\theta} - (1-\alpha)^2\|\theta\|_K^2 \le 0$$

for all $\tilde{\theta} \in \tilde{\Theta}^{\mathrm{BLS}}$ where $1 - (1-\alpha)^2 = \alpha(2-\alpha)$. Using the Cauchy-Schwarz inequality $-\theta^\top K\tilde{\theta} \le \|\theta\|_K\|\tilde{\theta}\|_K$, we obtain the following conservative condition

$$\alpha(2-\alpha)\|\tilde{\theta}\|_K^2 + 2(1-\alpha)^2\|\theta\|_K\|\tilde{\theta}\|_K - (1-\alpha)^2\|\theta\|_K^2 \le 0$$

for all $\tilde{\theta} \in \tilde{\Theta}^{\mathrm{BLS}}$. Since this quadratic equation is convex $\alpha(2-\alpha) > 0$, it is positive between its roots, which can be obtain from the quadratic formula

$$-\frac{1-\alpha}{\alpha}\|\theta\|_K \le \|\tilde{\theta}\|_K \le \frac{1-\alpha}{2-\alpha}\|\theta\|_K$$

for all $\tilde{\theta} \in \tilde{\Theta}^{\mathrm{BLS}}$. Since $\alpha \in (0,1)$ and $\|\tilde{\theta}\|_K \ge 0$, the lower-bound is redundant. Thus, $\tilde{\Theta}^{\mathrm{BLS}} \subseteq \tilde{\Theta}^{\max}$ if $\|\tilde{\theta}\|_K \le \frac{1-\alpha}{2-\alpha}\|\theta\|_K$ for all $\tilde{\theta} \in \tilde{\Theta}^{\mathrm{BLS}}$. For $\tilde{\theta} \in \tilde{\Theta}^{\mathrm{BLS}}$, the norm $\|\tilde{\theta}\|_K$ satisfies

$$\|\tilde{\theta}\|_K = \|K^{\frac{1}{2}}\tilde{\theta}\| \le \|K^{\frac{1}{2}}\|\|\tilde{\theta}\| = \|K^{\frac{1}{2}}\|\|\Lambda\Lambda^{-1}\tilde{\theta}\|$$
$$\le \|K^{\frac{1}{2}}\|\|\Lambda\|\|\Lambda^{-1}\tilde{\theta}\| \le \|K^{\frac{1}{2}}\|\|\Lambda\|$$

where $\|\Lambda^{-1}\tilde{\theta}\| \leq 1$ by the definition (18) of the set $\tilde{\Theta}^{\text{BLS}}$. Thus,

$$\tilde{\Theta}^{\text{BLS}} \subseteq \tilde{\Theta}^{\text{max}} \quad \text{if} \quad \|K^{\frac{1}{2}}\|\|\Lambda\| \leq \frac{1-\alpha}{2-\alpha}\|\theta\|_K$$

. Or equivalently, $\tilde{\Theta}^{\text{BLS}} \subseteq \tilde{\Theta}^{\text{max}}$ when the information matrix (4a) is sufficiently large

$$\Lambda^{-1} \succeq \frac{2-\alpha}{1-\alpha}\frac{\|K^{\frac{1}{2}}\|}{\|\theta\|_K}I \tag{24}$$

where $\|\theta\|_K = \|\nabla\mathcal{J}\|_K \neq 0$ by the hypothesis of this lemma. In other words, the ESC controller leaves the exploration mode when there is enough (24) information to reliably estimate the gradient. Next, we prove that the information matrix (4a) is sufficiently large (24) after a finite period of time $T < \infty$ in the exploration mode and non-zero dither amplitude $\delta > 0$. Consider the following two conditions:

1. The weightings (4c) are sufficiently large

$$w_k \geq \frac{2-\alpha}{1-\alpha}\frac{\|K^{\frac{1}{2}}\|}{\|\theta\|_K}\frac{1}{\rho\|\Delta y_k\|^2} \tag{25}$$

2. The output transients Δy_k are persistently exciting (PE)

$$\frac{1}{N}\sum_{k=t-N}^{t}\frac{\Delta y_k \Delta y_k^\top}{\|\Delta y_k\|^2} \succeq \rho I \tag{26}$$

where $\rho > 0$ quantifies the excitement of the outputs Δy_k. If conditions (25) and (26) hold then the information matrix (4a) is sufficiently large (24) to allow the controller (3) to enter the exploitation mode since

$$\begin{aligned}
\Lambda_t^{-1} &= \frac{1}{N}\sum_{k=t-1}^{t-N} w_k \Delta y_k \Delta y_k^\top \\
&\succeq \frac{1}{N}\sum_{k=t-1}^{t-N}\frac{2-\alpha}{1-\alpha}\frac{\|K^{\frac{1}{2}}\|}{\|\theta\|_K}\frac{\Delta y_k \Delta y_k^\top}{\rho\|\Delta y_k\|^2} \succeq \frac{2-\alpha}{1-\alpha}\frac{\|K^{\frac{1}{2}}\|}{\|\theta\|_K}I.
\end{aligned}$$

Thus, we next prove condition (25) holds after a finite-time $T_0 < \infty$ and non-zero dither amplitude $\delta > 0$. According to Corollary 3, the equilibrium state $\bar{x} = \pi(\bar{r})$ is ISS where the reference is constant $r_t = \bar{r}$ when the controller (3) is in the exploration mode. By the definition of ISS, we have

$$\|x_t - \bar{x}\| \leq \beta(\|x_0 - \bar{x}\|, t) + \gamma(\sup_t \|d_t\|)$$

where $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}$. Thus, for any initial condition, there exists a finite-time $T_0 < \infty$ and non-zero dither amplitude $\delta > 0$ such that

$$\|x_t - \bar{x}\| \leq \beta(\|x_0 - \bar{x}\|, T_0) + \gamma(\delta) \leq \frac{\rho}{\ell_g\|\tilde{H}\|}\frac{2-\alpha}{1-\alpha}\frac{\|K^{\frac{1}{2}}\|}{\|\theta\|_K}.$$

for all $t \geq T_0 \in \mathbb{N}$ and $\|d_t\| \leq \delta$ where ℓ_g is a Lipschitz bound on g. Since the plant (1) output map g is Lipschitz continuous, we have the following bound on the tracking error e_t

$$\|e_t\| = \|y_t - r_t\| = \|g(x_t) - g(\bar{x})\| \leq \ell_g\|x_t - \bar{x}\| \tag{27a}$$

$$\leq \frac{\rho}{\|\tilde{H}\|}\frac{2-\alpha}{1-\alpha}\frac{\|K^{\frac{1}{2}}\|}{\|\theta\|_K}. \tag{27b}$$

Likewise, we can obtain a bound on the output transients $y_k - y_{k-1}$

$$\|\Delta y_k\| = \|g(x_k) - g(\bar{x}) + g(\bar{x}) - g(x_{k-1})\| \tag{27c}$$

$$\leq \frac{2\rho}{\|\tilde{H}\|} \frac{2 - \alpha}{1 - \alpha} \frac{\|K^{\frac{1}{2}}\|}{\|\theta\|_K}.$$

Substituting the bounds (27) into the weightings (4c) produces

$$w_k^{-1} = \frac{1}{2} \|\Delta y_k\| \|\Delta y_k\|_{\tilde{H}} \left(\|e_t\|_{\tilde{H}} + \frac{1}{2} \|\Delta y_k\|_{\tilde{H}} \right)$$

$$= \frac{2 - \alpha}{1 - \alpha} \frac{\|K^{\frac{1}{2}}\|}{\|\theta\|_K} \rho \|\Delta y_k\|^2$$

Thus, condition (25) holds after finite-time $T_0 < \infty$ and for a non-zero dither amplitude $\delta > 0$.

Next, we prove condition (26) holds for a PE dither d_t. Since the plant (1) is controllable, a PE dither d_t will produce a PE state x_t after a finite-period n_c where n_c is the controllability index. Likewise, since the plant (1) is observable, the PE state will produce a PE output sequence Δy_k after a finite-period n_o where n_o is the observability index. Thus, condition (26) holds for some ρ after a finite period $n_c + n_o$. Note that the PE parameter ρ does not depend on the dither amplitude δ since the data $\{\Delta y_k\}_{k=t+n_c+n_o}^{t+n_c+n_o+N}$ is normalized in (26).

Thus, we have proven that conditions (25) and (26) hold after a finite period $T = \max\{T_0, n_c + n_o\}$ for a PE and sufficiently small $\|d_t\| \leq \delta$ dither d_t. This means that the information matrix (4a) is sufficiently large (24). Therefore, $\tilde{\Theta}^{\mathrm{BLS}} \subseteq \tilde{\Theta}^{\max}$ allowing the controller (3) to reenter the exploitation mode. $\square$

Lemma 2 shows that PE data $\{\mathcal{J}_k, y_k\}_{k=t-N}^{t}$ is necessary, but not sufficient to accurately estimate the cost gradient $\nabla \mathcal{J}(r_t)$. For an accurate estimates, the data must also be sufficiently local $y_t \approx r_t$ and sufficiently close to equilibrium $y_t \approx y_{t-1}$. For instance, data $\{\mathcal{J}_k, y_k\}_{k=t-N}^{t}$ collected far $\|y_k - r_t\| \gg 0$ from the set-point cannot be used to accurately estimate the gradient $\nabla \mathcal{J}(r_t)$ at r_t since the cost is nonlinear $H \preceq \nabla^2 \mathcal{J} \preceq \bar{H}$. Other ESC controllers (e.g. [3, 8]) address this issue by reducing the bandwidth of the ESC controller. In contrast, our ESC controller only reduces the bandwidth when it detects $\alpha_t^\star = 0$ that the gradient cannot be reliable estimated. Thus, we say our ESC controller has an adaptive separation of time-scales since the adaptive step-size (5) throttles the ESC controller (3) to allow the plant (1) settle providing better data for the gradient estimator (4).

The final result necessary for the proof of Theorem 1 connects the game-theoretic step-size (13) used in Corollary 3 with the step-size (5) used by the ESC controller (3).

Corollary 4. *Let Assumptions 1 and 2 hold. For the estimation error set* (18), *the optimal step-size* (13) *has the closed-form* (5).

Proof. For the estimation error set (18), the optimization problem (13) used to select the adaptive step-size $\alpha_t^\star$ can be reformulated as

$$\max_{\tilde{\theta} \in \tilde{\Theta}_t^{\mathrm{BLS}}} \tilde{\theta}^\top K \hat{\theta} = \max_{\|z\|^2 \leq 1} z^\top \Lambda_t^{\frac{1}{2}} K \hat{\theta}_t \tag{28}$$

where $z = \Lambda_t^{-1/2} \tilde{\theta}$ is a change-of-variables. The optimization problem (28) has a closed-form solution, namely

$$z^\star = \pm \Lambda_t^{\frac{1}{2}} K \hat{\theta}_t / \|\Lambda_t^{\frac{1}{2}} K \hat{\theta}_t\|$$

i.e. the unit vector $\|z^\star\| \leq 1$ aligned with the cost $\Lambda_t^{\frac{1}{2}} K \hat{\theta}_t$. Substituting $z^\star$ into (13) yields (5). $\square$

Corollary 4 shows that the step-size (5) used by the ESC controller (3) is closed-form solution of the game-theoretic optimal step-size (13) for the particular bounds (18) on the gradient estimation errors of the BLS estimator (4).

Finally, we can prove Theorem 1.

Proof of Theorem 1. Since the ESC controller (3) switches between the exploration and exploitation modes, we will use switched systems theory to prove stability.

Let t_i for $i \in \mathbb{N}$ denote the time-indices where the ESC controller (3) switches modes. Without loss of generality, assume that the system is in the exploration mode at even time-indices t_{2k} for $k \in \mathbb{N}$. From Corollaries 2 and 3, the common Lyapunov function (7) satisfies

$$V(\tilde{x}_{2k+2}, \tilde{r}_{2k+2}) - V(\tilde{x}_{2k}, \tilde{r}_{2k}) \tag{29}$$
$$\leq -q(\|\tilde{x}_{2k}, \tilde{r}_{2k}\|) + \sigma_s(\|d_{2k}, \ldots, d_{2k+1}\|)$$

where $\tilde{x}_{2k} = \tilde{x}_{t_{2k}}$ and $\tilde{r}_{2k} = \tilde{r}_{t_{2k}}$ and q is obtained by summing q_x and q_r along the closed-loop trajectories

$$q(\|\tilde{x}_{2k}, \tilde{r}_{2k}\|) = \sum_{t=t_{2k}}^{t_{2k+1}} q_x^1(\|\tilde{x}_t\|) + \sum_{t=t_{2k+1}}^{t_{2k+2}} q_x^2(\|\tilde{x}_t\|) + q_r^2(\|\tilde{r}_t\|).$$

where q is a class-$\mathcal{K}_\infty$ function of the states $(\tilde{x}_{2k}, \tilde{r}_{2k})$ at the $2k$ switching instance t_{2k}. Likewise, the summation

$$\sigma_s = \sum_{t=t_{2k}}^{t_{2k+2}} \sigma(\|d_t\|)$$

is a class-$\mathcal{K}$ function of the dither d_t between times t_{2k} and t_{2k+2}. Thus, by (29) and Proposition 2.3 in [23], there exists $\beta \in \mathcal{K}_\infty$ and $\gamma \in \mathcal{K}$ such that

$$\left\| \tilde{x}_{2k}, \tilde{r}_{2k} \right\| \leq \beta \left(\left\| \tilde{x}_0, \tilde{r}_0 \right\|, 2k \right) + \gamma(\delta)$$

Therefore, the state (x, r) of the closed-loop system (1) and (3)-(5) converges to a neighborhood of the optimal equilibrium (2) as the *switching index* k goes to infinity $k \to \infty$. Thus, we need to prove that switching index goes to infinity $k \to \infty$ as time goes to infinity $t \to \infty$ i.e. we do not become trapped in the exploration mode.

According to Lemma 2, if $k \not\to \infty$ then $t \not\to \infty$ since $t \leq \sum_{k=0}^{\sup k < \infty} T_k < \infty$ where $T_k < \infty$. Thus, by the contrapositive $k \to \infty$ as $t \to \infty$. Therefore, the optimal equilibrium (2) of the closed-loop system (1) and (3)-(5) is ISS. $\qquad\square$

4 Numerical Examples

In this section, we demonstrate our ESC controller through a series of numerical examples.

4.1 Illustrative Example

In this section, we demonstrate our ESC controller for a simple linear system with an unknown quadratic cost function. The purpose of this example is to illustrate our ESC controller (3)-(5) using classical control theory.

The plant (1) is an under-damped second-order linear system

$$\ddot{y} + 2\zeta\omega_n\dot{y} + \omega_n^2 y = \omega_n^2 u \tag{30}$$

where $\zeta = 0.1$ and $\omega_n = 1.0$. The cost $\mathcal{J}$ is a quadratic

$$\mathcal{J}(y) = \tfrac{1}{2}H(y - y^\star)^2 \tag{31}$$

where $y^\star$ is the optimal and $H \in \mathbb{R}$ is the Hessian.

Figure 1: The feedback-loop created by descending the gradient $H(y - y^\star)$ of the quadratic cost function (31).

First, we consider the ideal case where the gradient $\nabla \mathcal{J}(y) = H(y - y^\star)$ has been perfectly estimated $\tilde{\theta}_t = 0$. Since the gradient $\nabla \mathcal{J}(y) = H(y - y^\star)$ of the quadratic cost (31) is a linear function of the output y, we obtain the feedback-loop shown in Figure 1. Even though the gradient is perfectly estimated $\tilde{\theta}_t = 0$ and the plant (30) is open-loop stable, the integral-action of the ESC controller can destabilize the closed-loop system, as shown by the root-locus in Figure 2a. In particular, the Newton-step controller gain $\alpha_t K = H^{-1}$, which provides one-step convergence to the optimal for static optimization, destabilizes this dynamic optimization, as shown by root-locus in Figure 2a. In contrast, the proposed controller gain (6) provides closed-loop stability for perfect gradient estimates.


(a) Cost

(b) Step-Size

Figure 2: Root-locus plots demonstrating that the integral-action of the ESC controller (3) can destabilize the linear plant (30). Note that these are discrete-time root-locus plots.

16

Unfortunately, our instability issues re-emerge when we consider imperfect gradient estimates $\hat{\theta}_t \neq \nabla \mathcal{J}_t$. In particular, the worst-case gradient estimation error

$$\tilde{\theta}_t^\star = \frac{1-\alpha}{\alpha} \nabla \mathcal{J}_t \in \tilde{\Theta}_t = \tilde{\Theta}_t^{\max}$$

amplifies the feedback caused by the gradient

$$\hat{\theta}_t = \nabla \mathcal{J}_t + \tilde{\theta}_t = \frac{1}{\alpha} \nabla \mathcal{J}_t = \frac{H}{\alpha}(y - y^\star)$$

where $\alpha \ll 1$. This increases the loop-gain, leading to instability, as shown by the root-locus in Figure 2b. Fortunately, our adaptive step-size (5) will compensate $\alpha_t^\star = \alpha$ for the expansion $\frac{1}{\alpha}$ of the loop-gain to restore stability, as shown by the root-locus in Figure 2b.

(a) Cost

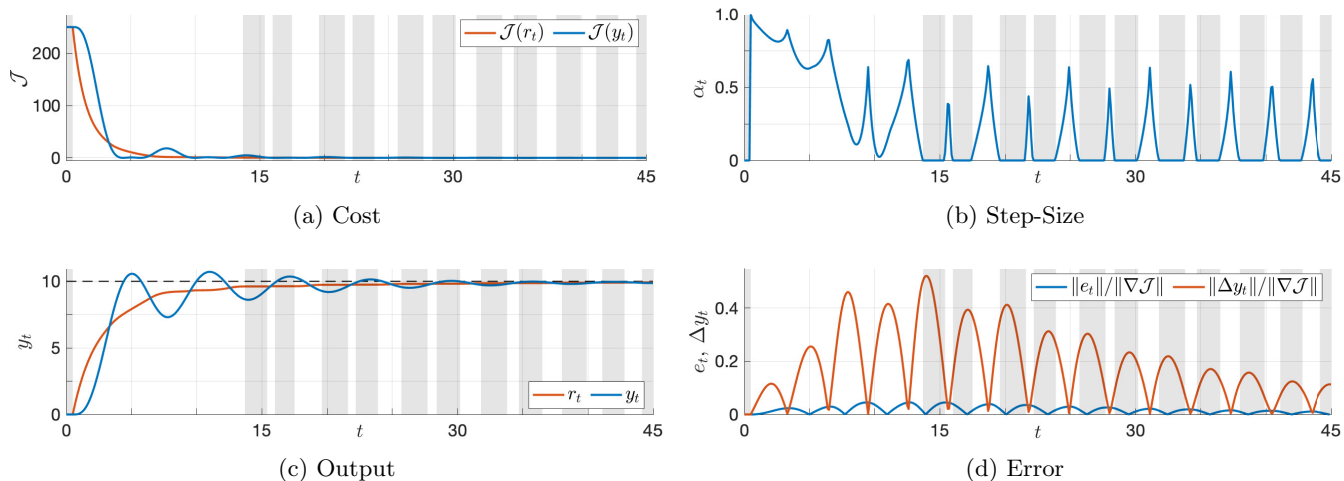(b) Step-Size

(c) Output

(d) Error

Figure 3: Closed-loop (30) and (3)-(5) simulation results. Shaded regions indicate when the ESC controller is in the exploration mode.

Finally, we demonstrate our ESC controller for the plant (30) and cost (31). The BLS estimator (4) has a batch horizon $N = 5$ and bounds $\underline{H} = 0$ and $\bar{H} = 10$ on the actual gradient $H = 5$. A dither $d_t = 0.001 \sin(t)$ was used to provide persistency of excitation.

Simulations for the closed-loop system (30) and (3)-(5) are shown in Figure 3. As shown in Figure 3c, the plant output converges $y_t \to y^\star$ to the optimal $y^\star = 10$. Since the plant (30) is under-damped, the output y_t oscillates and the measured cost $\mathcal{J}(y_t)$ converges non-monotonically to the optimal value. However, the set-point cost $\mathcal{J}(r_t)$ is monotonically decreasing as shown in Figure 3a. Figure 3b shows the step-size (5). Initially, the step-size (5) is large since the gradient $\nabla \mathcal{J} = H(y - y^\star)$ is large far from the optimal $|y - y^\star| \gg 0$. Thus, an accurate gradient estimate is not required to confidently descend. When the gradient becomes small, the controller often enters the exploration mode, indicated by the shaded regions in Figure 3. As shown in Figure 3d, the periods when the ESC controller is in the exploration mode $\alpha_t = 0$ correspond to periods when the plant is far from equilibrium $\|\Delta y_t\|/\|\nabla \mathcal{J}\| \gg 0$.

4.2 Practical Example: Drone Leak Inspection

In this section, we apply our ESC controller (3)-(5) to the problem of an autonomous drone searching for the source of an airborne pollutant leak.

The plant dynamics (1) model the closed-loop planar motion of a quadrotor drone. We use a standard model of the quadrotor dynamics e.g. [25]. For simplicity, we only consider the movement of the drone

in the plane i.e. the vertical position and orientation dynamics are ignored. The quadrotor is equip with GPS that measures its planar location $y \in \mathbb{R}^2$ and an integrated controller that moves the drone to a commanded location $y_t \to r \in \mathbb{R}^2$. Thus, the plant satisfies Assumption 1.

The objective of the ESC controller (3)-(5) is to move the drone to the source of a pollutant leak. The cost function $\mathcal{J}(y)$ optimized by the ESC controller (3) is the location y dependent measured concentration of pollutant in the air. Since our ESC controller minimize the cost function, we will consider the negative pollutant concentration. The negative pollutant concentration is modeled using a Gaussian plume model [26]

$$\mathcal{J}(y) = -\frac{1}{\sqrt{2\pi\sigma}} \exp\left(-\tfrac{1}{2}(y - y^\star)^\top \Sigma^\dagger (y - y^\star)\right) \tag{32a}$$

where $y^\star = [200, 100]^\top$ meters is the planar location of the leak and $\Sigma^\dagger(y - y^\star)$ is the pseudo-inverse of the covariance of pollutant concentration [26]

$$\Sigma(y - y^\star) = \begin{cases} \sigma^2(I - dd^\top) & \text{if } d^\top(y - y^\star) \geq 0 \\ \sigma_0^2 I & \text{otherwise} \end{cases} \tag{32b}$$

where $v = 10$ meters/second is the wind velocity (about 20 knots) and $d = [\cos(-\pi/4), \sin(-\pi/4)]^\top$ is the wind direction. The cost (32) says that the pollutant concentration has a Gaussian distribution in the cross-wind direction $(I - dd^\top)(y - y^\star)$. Note that the matrix $(I - dd^\top) = (I - dd^\top)^2$ is idempotent. The covariance $\sigma = \sigma_0 + d^\top(y - y^\star)/2$ of this Gaussian grows linearly with the distance $d^\top(y - y^\star)$ along the wind-direction d from the source $y^\star$. In the anti-wind direction $d^\top(y - y^\star) < 0$, the covariance is constant $\Sigma = \sigma_0^2 I$. Note that although the cost (32) is not convex, it locally satisfy Assumption 2.

Our discrete-time ESC controller (3)-(5) is executed at a rate of 20 Hertz. The BLS estimator (4) estimates the gradient $\nabla \mathcal{J}$ from the past 1 second of data, thus $N = 20$. The estimator uses the bounds $\underline{H} = -0.3I$ and $\bar{H} = 0.15I$ on the curvature $\nabla^2 \mathcal{J}$, which is approximately $5\times$ the actual curvature bounds. The controller gain (6) is $K = I$. The dither $d_t \sim \mathcal{N}(0, 1)$ is a normal distributed random variable with covariance of 1 meter.

Closed-loop simulation results are shown in Figures 4 and 5. Figure 4 shows the pollutant concentration and the path of the drone. The drone starts outside of the pollutant plume and moves perpendicular to the wind-direction into the plume stream. Once the drone enters the pollutant stream, is proceed against the wind direction to the source of the pollutant.

Figure 5a shows the location set-point r_t and concentration $\mathcal{J}(y_t)$ measured at y_t as a function of time t. The drone converges $r_t \to r^\star$ to the location $r^\star = y^\star$ of the pollutant leak as shown by the dashed black lines in Figure 5a. Likewise, the measured pollutant concentration $\mathcal{J}(y_t)$ converges to the maximum. The step-size (5) is shown in Figure 5b. Since the step-size (5) is zero $\alpha_t = 0$ approximately 90% of the time, we only plot it for the time-instance t when it is non-zero $\alpha_t \neq 0$. Since the ESC controller runs at 20 Hertz, the step-size is non-zero $\alpha_t \neq 0$ on average 12 times per minute, meaning that the estimated location r_t of the leak source is persistently and frequently updated.

4.3 Benchmark Examples

In this section, we compare our ESC controller with existing methods using three benchmark examples from the literature.

Figure 4: Pollutant concentration and drone path. Red indicates high concentration while blue indicates low concentration. Pollutant is blown by the wind to form a plume. The ESC controller drives the drone into the plume and follows the plume of the source of the leak.


(a) Cost

(b) Step-Size

Figure 5: (a) Reference location r_t for the drone and the resulting pollutant concentration $\mathcal{J}(y_t)$ over the 15 minute simulation. (b) Step-size α_t used by the ESC controller and the norms of the tracking error $e_t = r_t - y_t$ and drone transients $\Delta y_t = y_t - y_{t-1}$.

4.3.1 1-D Benchmark

In this section, we demonstrate our ESC controller for the 1 state benchmark example from [16]. The plant dynamics are

$$\dot{x} = -x + u \tag{33a}$$

$$y = x \tag{33b}$$

The plant (33) is a stable linear system and therefore satisfies Assumption 1. The unknown cost function is

$$\mathcal{J}(y) = 3 - \frac{1}{\sqrt{1 + (y - 2)^2}}. \tag{34}$$

Although the cost (34) is non-convex, it locally satisfies Assumption 2. Our ESC controller (3) used the gain $K = 0.5$ and a sample rate of 10 Hertz. The BLS estimator (4) had an estimation horizon of $N = 5$

and bounds $\underline{H} = -2$ and $\bar{H} = 2$ on the curvature $\nabla^2 \mathcal{J}$ of the cost. The dither $d_t = 0.001 \sin(t)$ was used to provide persistency of excitation.


(a) Cost

(b) Step-Size

Figure 6: A comparison of the proposed ESC controller with an ESC controller from literature.

Simulation results are shown in Figure 6. Figure 6a compares the cost $\mathcal{J}(y_t)$ profiles of our ESC controller with the ESC controller from [16]. This figure shows that our ESC controller converges to the optimal in roughly half the time as the existing controller.

4.3.2 2-D Benchmark

In this section, we demonstrate our ESC controller for the 2 state benchmark example from [20]. The plant dynamics are

$$\dot{x} = R(x)u + w \tag{35a}$$
$$y = x \tag{35b}$$

where $R(x) \in \mathbb{R}^{2 \times 2}$ is a planar rotation matrix with angle $x_1 + x_2$ and $w(t) = [\sin(2t), \cos(t)]$ is a periodic disturbance. This nonlinear plant (35) does not satisfy Assumption 1 since it is not ISS. Indeed, it is only *marginally stable* for $u = w = 0$ and has no equilibrium states $\dot{x} = 0$ for $u \neq 0$ or $w \neq 0$. Thus, we pre-stabilize the system using the controller

$$u = -R(x)^\top \big(F(x - r) - w\big).$$

where the matrix $F = -10I$ has Hurwitz eigenvalues so that the output will track $y \to r$ the reference $r = \bar{r}$. To make the problem more challenging and preserve the nonlinearity, we simulate the plant (35) in continuous-time with the controller updated in discrete-time, i.e., we apply a zero-order hold for the control input $u(t) = u(t_k)$ for $t \in [t_k, t_{k+1})$ which is computed for states $x(t_k)$ and disturbances $w(t_k)$ sampled as discrete-times t_k where $\Delta t = 50$ milliseconds.

The unknown cost function is

$$\mathcal{J}(y) = \big\|y - 1\big\|^2 + 2018. \tag{36}$$

This strictly convex quadratic cost satisfies our Assumption 2.

The ESC controller (3) used gain $K = 0.5I$ and a sample-rate of 20 Hertz. The BLS estimator (4) had an estimation horizon of $N = 5$ and bounds $\underline{H} = 0I$ and $\bar{H} = 10I$ on the curvature $\nabla^2 \mathcal{J}$ of the cost. No dither was used since the periodic disturbance $w(t)$ already provide persistency of excitation. Between sample periods $\Delta t = 0.05$, the nonlinear plant (35) was simulated using MATLAB's `ode45` solver.

Simulation results are shown in Figure 7. Figure 7a shows that our ESC controller has comparable performance to the existing controller from [20]. This benchmark example demonstrates how the tracking errors $e_t = r_t - y_t$ affect our adaptive step-size (5). Due to the rotation matrix in the dynamics (35), the plant takes looping paths between the reference set-points r_t. This produces highly exciting, but highly

(a) Cost

(b) Step-Size

Figure 7: A comparison of the proposed ESC controller with an existing ESC controller from literature.

non-local $\|r_t - y_t\| \gg 0$ data $\{\mathcal{J}_t, y_t\}$, which leads to poor estimates of the gradient $\nabla \mathcal{J}(r_t)$ at the set-point r_t. As a result, the step-size (5) is almost always zero $\alpha_t = 0$, allowing the plant (35) to settle $y_t \approx r_t$ near the set-point r_t before trusting the estimated gradient. Indeed, the step-size is non-zero $\alpha_t \neq 0$ at only 9 of the 501 simulated time instances, as shown in Figure 7b. Nonetheless, our ESC controller converged to the optimal with a comparable convergence rate to the specialized ESC controller from [20].

4.3.3 3-D Benchmark

In this section, we demonstrate our ESC controller for the 3 state benchmark example from [13]. The plant dynamics are

$$\dot{x}_1 = -x_1 + u_2^2 \tag{37a}$$
$$\dot{x}_2 = -x_2 + u_1 \tag{37b}$$
$$\dot{x}_3 = -x_3 + u_2 x_2. \tag{37c}$$

Although the plant (1) does not satisfy our asymptotic tracking assumption, this can be rectified by inverting the steady-state map of the plant using the transformation

$$u_1 = r_1/(1 + \sqrt{r_2}) \tag{38a}$$
$$u_2 = \sqrt{r_2}. \tag{38b}$$

The plant (37) has an implicit constraint $r_2 \geq 0$, which we enforce by setting $r_2 = 0$ if $r_2 < 0$. This plant (37) is only locally Lipschitz continuous. The unknown cost function is

$$\mathcal{J}(y) = (x_2 + x_3)^2 + 2(x_1 + x_2 - u_1) = y_1^2 + 2y_2 \tag{39}$$

where $y_1 = x_2 + x_3$ and $y_2 = x_1 + x_2 - u_1$ are the measured outputs. Note that the cost (39) is convex, but not strictly convex. Nonetheless, it satisfies Assumption 2.

For the ESC controller (3) design, the plant (37) was converted to discrete-time using the forward Euler method with a sample-time of $\Delta t = 0.25$. The gain G and Lyapunov matrix P were computed using parametric linear matrix inequalities [27] with $u_2 \in [-5, 5]$ as the parameter. The BLS estimator (4) had an estimation horizon of $N = 5$ and bounds $\underline{H} = 0I$ and $\bar{H} = 10I$ on the curvature $\nabla^2 \mathcal{J}$ of the cost. The dither $d_t = 0.001[\sin(t), \sin(2t)]^\top$ was used to provide persistency of excitation.

Simulation results are shown in Figure 8. Between sample periods $\Delta t = 0.25$, the nonlinear plant (37) was simulated using MATLAB's `ode45` solver. Figure 8a shows that our ESC controller converged to the optimal equilibrium in approximately 30 seconds, which is approximate $40\times$ faster than the existing controller. Simulation results for the existing controller are not shown due to the disparity in time-scales. Note that the cost $\mathcal{J}(y_t)$ can temporarily drop below the optimal equilibrium cost since $y_2 = x_1 + x_2 - u_1 = x_1 - \dot{x}_2$ depends on the state velocity $\dot{x}_2$ which is zero $\dot{x}_2 = 0$ at equilibrium.

21

(a) Cost

(b) Step-Size

Figure 8: A comparison of the proposed ESC controller with an existing ESC controller from literature.

Again, the step-size (5) is almost always zero α_t as shown in Figure 7b. The step-size is non-zero $\alpha_t \neq 0$ for only 4 of the simulated 120 time instances. In contrast to the previous benchmark example, in this example the mostly zero step-size $\alpha_t = 0$ is due to the state velocity $\dot{x}_2$ appearing in the cost. The zero step-size $\alpha_t = 0$ allows the plant (37) to settle near an equilibrium $y_t \approx y_{t-1}$ before exploiting the estimated gradient.

5 Conclusions

This paper presented an ESC controller (3) with an adaptive step-size (5) that adjusts the aggressiveness of the controller based on the quality of the gradient estimate (4). We proved that the BLS estimator (4) with our novel weighting (4c) produced bounded (18) gradient estimation errors. The adaptive step-size (5) maximizes the decrease of the Lyapunov function (7) for the worst-case estimation error (18) in the exploitation mode. In the exploration mode, the controller allows the plant to settle improving the gradient estimate. Since the controller (3) interminably re-enters the exploitation mode, we were able to prove that the optimal equilibrium (2) is ISS for the closed-loop system (1) and (3).

References

[1] M. LeBlanc, "Sur l'electrifaction des chemins de fer au moyen de courantsalternatifs de frequence elevee," *Rev. Gen. l'Electr*, vol. 2, pp. 275–277, 1922.

[2] K. T. Atta and M. Guay, "Adaptive amplitude fast proportional integral phasor extremum seeking control for a class of nonlinear system," *Journal of Process Control*, vol. 83, pp. 147 – 154, 2019.

[3] M. Guay and D. J. Burns, "A proportional integral extremum-seeking control approach for discrete-time nonlinear systems," in *2015 54th IEEE Conference on Decision and Control (CDC)*, 2015, pp. 6948–6953.

[4] M. Guay and D. Dochain, "A time-varying extremum-seeking control approach," *Automatica*, vol. 51, pp. 356–363, 2015.

[5] C. Danielson, S. Lacy, B. Hindman, a. H. P. Collier, and R. Moser, "Extremum seeking control for simultaneous beam steering and wavefront correction," in *American Control Conference*, 2006.

[6] E. Biyik and M. Arcak, "Gradient climbing in formation via extremum seeking and passivity-based coordination rules," in *2007 46th IEEE Conference on Decision and Control*, 2007, pp. 3133–3138.

[7] A. Chakrabarty, C. Danielson, S. Di Cairano, and A. Raghunathan, "Active learning for estimating reachable sets for systems with unknown dynamics," *IEEE Transactions on Cybernetics*, pp. 1–12, 2020.

[8] K. Ariyur and M. Krstic, *Real-Time Optimization by Extremum-Seeking Control.* Wiley, 2003.

[9] M. Benosman, *Learning-Based Adaptive Control: An Extremum Seeking Approach - Theory and Applications.* Butterworth-Heinemann, 2017.

[10] C. Yin, S. Dadras, X. Huang, Y. Chen, and S. Zhong, "Optimizing energy consumption for lighting control system via multivariate extremum seeking control with diminishing dither signal," *IEEE Transactions on Automation Science and Engineering*, vol. 16, no. 4, pp. 1848–1859, 2019.

[11] L. Wang, S. Chen, and K. Ma, "On stability and application of extremum seeking control without steady-state oscillation," *Automatica*, vol. 68, pp. 18–26, 2016.

[12] Z. He, S. Chen, Z. Sun, L. Wang, and K. Ma, "Twisting sliding mode extremum seeking control without steady-state oscillation," *International Journal of Control*, pp. 1–11, 2019.

[13] M. Haring and T. A. Johansen, "Asymptotic stability of perturbation-based extremum-seeking control for nonlinear plants," *IEEE Transactions on Automatic Control*, vol. 62, no. 5, pp. 2302–2317, 2016.

[14] A. Scheinker and M. Krstić, "Minimum-seeking for CLFs: Universal semiglobally stabilizing feedback under unknown control directions," *IEEE Trans. Automat. Contr.*, vol. 58, no. 5, pp. 1107–1122, 2013.

[15] A. Ghaffari, M. Krstić, and S. Seshagiri, "Power optimization for photovoltaic microconverters using multivariable newton-based extremum seeking," *IEEE Trans. Control Syst. Technol.*, vol. 22, no. 6, pp. 2141–2149, 2014.

[16] B. G. Hunnekens, M. A. Haring, N. Van De Wouw, and H. Nijmeijer, "A dither-free extremum-seeking control approach using 1st-order least-squares fits for gradient estimation," *Proc. IEEE Conf. Decis. Control*, vol. 2015-Febru, no. February, pp. 2679–2684, 2014.

[17] J. I. Poveda and A. R. Teel, "A robust event-triggered approach for fast sampled-data extremization and learning," *IEEE Trans. Automat. Contr.*, vol. 62, no. 10, pp. 4949–4964, 2017.

[18] T. I. Salsbury, J. M. House, and C. F. Alcala, "Self-perturbing extremum-seeking controller with adaptive gain," *Control Engineering Practice*, vol. 101, p. 104456, 2020.

[19] V. Grushkovskaya, A. Zuyev, and C. Ebenbauer, "On a class of generating vector fields for the extremum seeking problem: Lie bracket approximation and stability properties," *Automatica*, vol. 94, pp. 151–160, 2018. [Online]. Available: https://doi.org/10.1016/j.automatica.2018.04.024

[20] R. Suttner, "Extremum seeking control with an adaptive dither signal," *Automatica*, vol. 101, pp. 214–222, 2019. [Online]. Available: https://doi.org/10.1016/j.automatica.2018.11.055

[21] Z. Jiang and Y. Wang, "A converse Lyapunov theorem for discrete-time systems with disturbances," *Syst. Control. Lett.*, vol. 45, pp. 49–58, 2002.

[22] C. M. Kellett and A. R. Teel, "Results on discrete-time control-Lyapunov functions," in *42nd IEEE International Conference on Decision and Control*, vol. 6, 2003, pp. 5961–5966 Vol.6.

[23] Z.-P. Jiang and Y. Wang, "Input-to-state stability for discrete-time nonlinear systems," *Automatica*, vol. 37, no. 6, pp. 857 – 869, 2001.

[24] Y. Xu, W. Sun, and L. Qi, "A feasible direction method for the semidefinite program with box constraints," *Applied Mathematics Letters*, vol. 24, no. 11, pp. 1874 – 1881, 2011.

[25] T. Bresciani, "Modeling, Identification, and Control of a Quadrotor Helicopter," Master's thesis, Lund University, 2008.

[26] M. R. Beychok, *Fundamentals of Stack Gas Dispersion*. M.R. Beychok, 1994.

[27] S. Boyd, L. Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. Society for Industrial and Applied Mathematics, 1994.