Self-Optimizing Traffic Light Control Using Hybrid Accelerated Extremum Seeking

Felipe Galarza-Jimenez, Jorge I. Poveda, Ronny Kutadinata, Lele Zhang, Emiliano Dall'Anese

***Abstract*— Motivated by the shallow concavity properties that emerge in certain response maps in the context of optimization problems in transportation systems, we study the stability properties of a class of hybrid accelerated extremum seeking (HAES) dynamics interconnected with dynamic plants in the loop. In particular, we establish suitable semi-global practical asymptotic stability properties for different classes of cost functions, as well as tuning conditions for the hybrid extremum seeking algorithm. Additionally, we implement the HAES to optimize the performance of a self-organizing traffic light system (SOTL) in a class of smart transportation systems. We show that the dynamic momentum mechanism incorporated by the HAES can significantly reduce the convergence time in the optimization process compared to the traditional extremum seeking algorithms based on gradient descent flows.**

I. Introduction

Technological advances in sensing, computation, and communication have enabled the deployment of novel automation technologies in networked transportation systems. As a prime example, traffic light control has emerged as a fundamental technology for the efficient operation of urban cities [1]. In these systems, the feedback control mechanism must be able to *leverage data* to bypass the lack of traffic models, and also to *rapidly* cope with changes in traffic conditions. Moreover, the controllers should be *adaptive* and *robust* with respect to unexpected disturbances, such as demand fluctuations, traffic accidents, etc [2].

On the other hand, practitioners may have access only to real-time measurements from sensing units and recorded data. These streams of data can be used to develop feedback-based *data-driven* algorithmic frameworks that can "learn to optimize" the traffic behavior in real-time. This practical setting motivates this paper, which studies the real-time optimization of traffic systems using efficient *extremum seeking control* (ESC) [3]–[6], a type of adaptive control suitable for the solution of self-optimization problems in dynamical systems. Along with the advantages of ESC, it has been observed in several applications -particularly those with cost functions having shallow convexity/concavity properties- that the method might exhibit slow rates of convergence. Such limitations are structural and induced by the type of optimization algorithm that is most commonly used in the literature of ESC: gradient descent flows. Indeed, as shown in [7], the rate of convergence of gradient flows can be of order $\mathcal{O}(1/t)$ in the family of smooth convex cost functions. In the optimization literature, this limitation has motivated several works on "accelerated" methods able to achieve rates of convergence of order $\mathcal{O}(1/t^2)$ for convex functions, and of order $\mathcal{O}(e^{-\sqrt{\kappa}t})$ for κ-strongly convex functions with $\kappa > 0$ [8], [9]. However, as shown in [10], such dynamics are incompatible with ESC due to their lack of uniformity in the convergence properties. To solve the incompatibility issue, a class of hybrid accelerated extremum seeking (HAES) dynamics was introduced in [11] for the optimization of *static maps* with unknown mathematical forms. These algorithms can achieve full acceleration properties in strongly convex functions and semi-acceleration properties in any smooth radially unbounded convex function. However, it has remained an open question whether these ES dynamics can also be used to optimize *dynamical systems* in the loop instead of static maps.

In this paper, we provide a positive answer to this question by showing that the HAES, interconnected with a stable plant generating a well-defined quasi-steady state cost function, has suitable semi-global practical asymptotic stability properties and preserves the acceleration bounds as the time-scale separation in the closed-loop system increases. Moreover, we show that the qualitative tuning guidelines developed for static optimization problems, and related to the frequency of the resets of the hybrid controller, are also extendable to the dynamic case.

After establishing suitable stability properties for the accelerated ESC with plants in the loop, we proceed to apply the method to solve an optimization problem in transportation systems, where the goal is to optimize in real-time the performance of *self-organizing traffic light systems* (SOTL). SOTL is a control strategy at the intersection level that utilizes a set of threshold-based rules. As discussed in [12], the performance of SOTL depends upon the controller thresholds, which impact the mean travel time in the transportation network. However, the mathematical form of the mapping between controller thresholds and travel time is not available, and therefore cannot be optimized offline. Nevertheless, in the case of SOTL, the work [2] showed that the concavity of the Macroscopic Fundamental Diagram (MFD) is inherited by the parameter-to-average speed map, and that it exhibits regions with shallow concavity, i.e., almost "flat" regions. Therefore, even if the controller has access to real-time streams of data for the purpose of online optimization, standard ESC algorithms based on gradient descent flows

F. Galarza-Jimenez, J. I. Poveda and E. Dall'Anese are with the Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO. Contact Email: jorge.poveda@colorado.edu

R. Kutadinata is with the Mechatronics Department, Deakin University, Melbourne, Australia.

L. Zhang is with the School of Mathematics and Statistics, University of Melbourne, Australia.

Research supported in part by NSF grant CRII - CNS: 1947613 and NSF grant CMMI 2044946.

are expected to be slow due to the shape of the response map of the system. We address this issue by implementing the HAES to optimize the performance of the SOTL.

II. Preliminaries

A. Notation

The set of (non-negative) real numbers is denoted as ($\mathbb{R}_{\geq 0}$) $\mathbb{R}$. The set of (non-negative) integers is denoted as ($\mathbb{Z}_{\geq 0}$) $\mathbb{Z}$. Given a compact set $\mathcal{A} \subset \mathbb{R}^n$, and a column vector $x \in \mathbb{R}^n$, we define $|x|_{\mathcal{A}} := \min_{y \in \mathcal{A}} |x - y|$. We use $\mathbb{B}$ to denote a closed unit ball of appropriate dimension, $\rho\mathbb{B}$ to denote a closed ball of radius $\rho > 0$, and $\mathcal{X} + \rho\mathbb{B}$ to denote the union of all sets obtained by taking a closed ball of radius ρ around each point in the set $\mathcal{X}$. In this paper, we use $\mathcal{R}_\kappa$ to denote a class of $2n \times 2n$ block matrices, with individual block component $\mathcal{R}_i = [0, \kappa_i; -\kappa_i, 0]$, where $\kappa_i > 0$, for $i \in \{1, 2, \ldots, n\}$. We also use $\mathbb{T}^n := \mathbb{S}^1 \times \mathbb{S}^1 \times \ldots \times \mathbb{S}^1$ to model the n^{th} Cartesian product of unit circles.

B. Hybrid Dynamical Systems

We model our algorithms as hybrid dynamical systems (HDS) [13]. In particular, we consider HDS with state $\zeta \in \mathbb{R}^n$, evolving according to:

$$\dot{\zeta} = F(\zeta), \qquad \zeta \in C, \tag{1a}$$
$$\zeta^+ = G(\zeta), \qquad \zeta \in D, \tag{1b}$$

where $\dot{\zeta}$ stands for the derivative of ζ with respect to time, and ζ^+ corresponds to the value of ζ after an instantaneous change. The mappings $F : \mathbb{R}^n \to \mathbb{R}^n$ and $G : \mathbb{R}^n \to \mathbb{R}^n$ are called the flow map and the jump map, respectively, and they describe the evolution of the system when ζ belongs to the flow set C, or/and the jump set D, respectively. System (1) is represented by the notation $\mathcal{H} := \{C, F, D, G\}$, where C, F, D and G comprise the *data* of $\mathcal{H}$[1].

For a HDS parametrized by a constant $\varepsilon \in \mathbb{R}_{>0}$, denoted as $\mathcal{H}_\varepsilon := \{C_\varepsilon, F_\varepsilon, D_\varepsilon, G_\varepsilon\}$, a compact set $\mathcal{A} \subset \mathbb{R}^n$ is said to be semi-globally practically asymptotically stable (SGPAS) as $\varepsilon \to 0^+$ if there exists a function $\beta \in \mathcal{KL}$[2] such that the following holds: For each pair $K > \nu > 0$ there exists an $\varepsilon^* \in \mathbb{R}_{>0}$ such that for each $\varepsilon \in (0, \varepsilon^*)$ each solution ζ of $\mathcal{H}_\varepsilon$ that satisfies $|\zeta(0,0)|_{\mathcal{A}} \leq K$ also satisfies $|\zeta(t,j)|_{\mathcal{A}} \leq \beta(|\zeta(0,0)|_{\mathcal{A}}, t + j) + \nu$, for all $(t,j) \in \text{dom}(\zeta)$. If β has exponential form, we use the acronym SGPES. If instead of K we consider $\mathbb{R}^n$, and make $\nu = 0$, the set $\mathcal{A}$ is said to be uniformly globally asymptotically stable (UGAS).

[1]Note that HDS of the form (1) are a generalization of purely continuous-time systems and purely discrete-time systems. Namely, continuous-time dynamical systems can be seen as a HDS of the form (1) with $D = \emptyset$, while discrete-time dynamical systems can be seen as a HDS of the form (1) with $C = \emptyset$. For a complete definition of solutions to systems of the form (1) we refer the reader to [13, Ch. 2].

[2]For a definition of functions of class $\mathcal{KL}$ we refer the reader to [13, Ch. 3.5].

III. Accelerating Extremum Seeking using Dynamic Momentum

ESC describes a family of adaptive controllers developed to solve model-free optimization problems in dynamic systems. In particular, ESC considers plants of the form

$$\dot{\theta} = f(\theta, u), \qquad y = h(\theta, u), \tag{2}$$

where, traditionally, it is assumed that f and h are sufficiently smooth, $\theta \in \mathbb{R}^p$, $u \in \mathbb{R}^n$, $y \in \mathbb{R}$ and that the following standard assumption holds.

Assumption 1: There exists a continuous function $\ell : \mathbb{R}^n \to \mathbb{R}^p$, such that for each compact set $K_u \subset \mathbb{R}^n$, the dynamical system

$$\dot{\theta} = f(\theta, u), \quad \dot{u} = 0, \quad (\theta, u) \in \mathbb{R}^p \times K_u \tag{3}$$

renders UGAS the set $M_{K_u} := \{(\theta^*, u^*) \in \mathbb{R}^p \times \mathbb{R}^n : \theta^* = \ell(u), u \in K_u\}$. □

Under Assumption 1, the cost function in ESC is defined as the following steady state input-to-output map, also called, the *response map* of (2):

$$\phi(u) := h(\ell(u), u), \tag{4}$$

which is assumed to satisfy the following:

Assumption 2: The function ϕ satisfies $\min_{u \in \mathbb{R}^n} \phi > -\infty$ as well as the following properties: 1) $\phi(\cdot)$ is continuously differentiable. 2) ϕ is radially unbounded. 3) ϕ is convex.□

Under Assumption 2, the set of minimizers of ϕ, denoted $\mathcal{A}$, is non-empty and compact. We define $\phi^* := \phi(\mathcal{A})$. To achieve exponential rates of convergence, the following assumption is usually made:

Assumption 3: The function ϕ continuously differentiable and strongly convex, i.e., there exist $\kappa > 0$, such that

$$\phi(u_1) - \phi(u_2) + \nabla\phi(u_2)^\top (u_2 - u_1) \geq \frac{\kappa}{2}|u_2 - u_1|^2, \tag{5}$$

for all $u_1, u_2 \in \mathbb{R}^n$. Moreover, $\nabla\phi$ is globally Lipschitz with constant $L > 0$. □

Under Assumption (3), the function ϕ is radially unbounded, and the minimizer of ϕ is unique, i.e., $\mathcal{A} := \{u^*\}$.

A. Gradient Descent-Based ESC

To minimize the cost function (4), the traditional gradient descent-based ESC, introduced in [3], and further studied in [14], takes the form

$$\dot{x} = -\varepsilon_o k \frac{2}{\varepsilon_a} \mu\left(\frac{2\pi\varepsilon_o \boldsymbol{\kappa} t}{\varepsilon_p}\right) y, \quad u = x + \varepsilon_a \mu\left(\frac{2\pi\varepsilon_o \boldsymbol{\kappa} t}{\varepsilon_p}\right), \tag{6}$$

which is interconnected in feedback with system (2). In (6), the constant $k > 0$ acts as a gain of the controller, the function $\mu : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ is a vector of element-wise sinusoidal signals with frequencies parameterized by the constants $\boldsymbol{\kappa} = [\kappa_1, \kappa_2, \ldots, \kappa_n]^\top$, which satisfy the following:

Assumption 4: The parameters κ_i are positive rational numbers, and they satisfy $\kappa_i \neq \kappa_j$ and $\kappa_i \neq 2\kappa_j$, for all $i \neq j$. □

In (6), parameters $(\varepsilon_o, \varepsilon_a, \varepsilon_p)$ are selected sufficiently small to induce three time scales in the closed-loop system: the fastest time scale is related to the dynamics of the plant (2), the medium-time scale corresponds to the oscillations of μ, and the slowest time scale corresponds to the variations of x. This intuitive behavior can be formalized using tools from averaging and singular perturbation theory; see [3]. In particular, using the change of time variable $s = \varepsilon_o t$, the closed-loop system can be written in the s-time scale as

$$\varepsilon_o \frac{d\theta}{ds} = f\left(\theta, x + \varepsilon_a \mu\left(\frac{2\pi \boldsymbol{\kappa} s}{\varepsilon_p}\right)\right), \tag{7a}$$

$$\frac{dx}{ds} = -k\frac{2}{\varepsilon_a}\mu\left(\frac{2\pi \boldsymbol{\kappa} s}{\varepsilon_p}\right) h\left(\theta, x + \varepsilon_a \mu\left(\frac{2\pi \boldsymbol{\kappa} s}{\varepsilon_p}\right)\right). \tag{7b}$$

System (7) is a singularly perturbed system, with ε_o acting as small parameter. As $\varepsilon_0 \to 0^+$, the behavior of x is mainly predicted by the so-called *reduced dynamics*, which assume that the plant dynamics (7a) is at equilibrium, i.e., $\theta = \ell\left(x + \varepsilon_a \mu\left(\frac{2\pi \boldsymbol{\kappa} s}{\varepsilon_p}\right)\right)$. Using the definition of (4), we obtain:

$$\frac{d\tilde{x}}{ds} = -k\frac{2}{\varepsilon_a}\mu\left(\frac{2\pi \boldsymbol{\kappa} s}{\varepsilon_p}\right)\phi\left(\tilde{x} + \varepsilon_a \mu\left(\frac{2\pi \boldsymbol{\kappa} s}{\varepsilon_p}\right)\right). \tag{8}$$

For small values of ε_p system (8) can be analyzed via averaging theory. Under Assumption 4, standard computations based on Taylor expansions (see [5]) reveal that the average system of (8) is given by

$$\frac{dz}{ds} = -k\nabla\phi(z) + \mathcal{O}(\varepsilon_a). \tag{9}$$

Since $\mathcal{O}(\varepsilon_a)$ is a small bounded disturbance (on compact sets), it can be treated as structural additive perturbation. Thus, the nominal dynamics of (9) is just a gradient flow

$$\frac{d\tilde{z}}{ds} = -k\nabla\phi(\tilde{z}). \tag{10}$$

Under Assumption 2, gradient flows of the form (10) render the set $\mathcal{A}$ UGAS. Therefore, solutions of (10) satisfy

$$|\tilde{z}(s)|_{\mathcal{A}} \leq \beta_{opt}(|\tilde{z}(0)|_{\mathcal{A}}, s), \tag{11}$$

for some $\beta_{opt} \in \mathcal{KL}$. Moreover, in the class of cost functions satisfying Assumption 2, gradient flows of the form (10) minimize ϕ at a worst-case rate of $O(1/s)$ [9]. If we further assume that ϕ satisfies condition (5) with a L-globally Lipschitz gradient, it can be shown that (11) holds with $\beta_{opt}(r, s) = \sqrt{L/\kappa}e^{-\kappa k s}r$, i.e., the convergence is of order $\mathcal{O}(e^{-k\kappa s})$. Standard robustness results for dynamical systems (e.g., [13, Prop. 6.34]) predict that, as $\varepsilon_a \to 0^+$, the solutions of (10) will be $\epsilon/3$-close (on compact time domains) to the solutions of the perturbed system (9). Moreover, standard results in averaging theory predict that as $\varepsilon_p \to 0^+$ the trajectories of (9) will be $\epsilon/3$-close (on compact time domains) to the solutions of system (8). Additionally, by standard results in singular perturbation theory, the trajectories $\tilde{x}$ of (8) will be $\epsilon/3$-close (on compact time domains) to the trajectories x of system (7). By combining these closeness properties with the $\mathcal{KL}$ bound (11), one can establish that for any $K > v > 0$ there exist $\varepsilon_a^* > 0$ such that for any $\varepsilon_a \in (0, \varepsilon_a^*)$ there exists $\varepsilon_p^* > 0$ such that for any $\varepsilon_p \in (0, \varepsilon_p^*)$ there exists $\varepsilon_o \in (0, \varepsilon_o^*)$ such that the trajectories of the original closed-loop system with $|\theta(0)|_{\ell(\mathcal{A})} \leq K$ and $|x(0)|_{\mathcal{A}} \leq K$ satisfy

$$|x(t)|_{\mathcal{A}} \leq \beta_{opt}(|x(0)|_{\mathcal{A}}, \varepsilon_o t) + \nu, \tag{12}$$

for all $t \geq 0$.

The above discussion highlights an important property in ESC: as the time-scale separation increases in the closed-loop system (7), the transient performance of the input is characterized by the nominal average system (10). Given that gradient flows can be extremely slow when $\kappa \ll 1$, it is natural to ask how to improve the convergence rate of the algorithm without destroying the stability properties of the system. Naturally, increasing the gain k is not a feasible approach when there is a plant in the loop, as in (7), since this would destroy the time-scale separation in the system. Similarly, Newton-like ESC [15] requires the computation and inversion -on average- of the Hessian matrix of ϕ, which must be positive definite, i.e., ϕ must satisfy (5). A similar limitation applies to finite-time or fixed-time ESCs [16]. On the other hand, a different approach is to consider an ESC algorithm with an average system that generates more "efficient" optimization directions compared to the standard gradient flow (10). Algorithms with this *geometric* property are term "accelerated", and they are the subject of the next section.

B. ESC with Dynamic Momentum: Instability

Motivated by Nesterov's accelerated optimization algorithm [17], momentum techniques have recently received significant attention in the literature of optimization and dynamical systems. In particular, in [8] and [9], the authors showed that the continuous-time version of Nesterov's algorithm is given by a second-order ODE of the form

$$\ddot{x} + \frac{c}{t}\dot{x} + k\nabla\phi(x) = 0, \quad c \geq 3, \tag{13}$$

which, under Assumption 2, minimizes ϕ at a rate of $\mathcal{O}(1/t^2)$ in the family of smooth convex functions. Thus, it is natural to consider system (13) as a candidate to substitute the dynamics (6) by the following ESC with dynamic momentum:

$$\dot{x} = \varepsilon_o q, \quad \dot{q} = -\frac{c\varepsilon_o q}{t} - \varepsilon_o k\frac{2}{\varepsilon_a}\mu\left(\frac{2\pi\varepsilon_o t}{\varepsilon_p}\right) y. \tag{14}$$

Unfortunately, in general, this system can become unstable under the slightest disturbance (e.g., noise) on the cost y, no matter how small are the parameters $(\varepsilon_o, \varepsilon_p, \varepsilon_a)$. This instability is a consequence of structural properties of the dynamics (13), studied in [10, Ex. 1] and [18]. Namely, in general, for these dynamics there does not exist a class $\mathcal{KL}$ function β_{opt} that satisfies (11). This issue can be addressed by incorporating into (14) mechanisms that prevent the damping from vanishing to zero, while further enhancing the transient performance of the dynamics. Such mechanisms can be designed using the framework of hybrid dynamical systems.

C. Robust Accelerated Hybrid ESC for Static Maps

To endow the dynamics (14) with suitable stability and robustness properties, we consider the incorporation of resets into the ESC. In particular, we introduce a new variable τ to model the evolution in time of the damping coefficient of (13), and a state μ, evolving on the tori $\mathbb{T}^n$, to model a vector of probing sinusoidal signals generated by a time-invariant dynamic oscillator. Using these additional auxiliary states, we consider a class of hybrid accelerated extremum seeking (HAES) dynamics for static maps (i.e., with $\varepsilon_o = 1$) with overall state $z = (x, p, \tau, \mu)$, continuous-time dynamics given by

$$\dot{x} = \frac{2}{\tau}(p - x), \quad \dot{p} = -\frac{4\tau k}{\varepsilon_a}\tilde{\mu}\phi(x + \varepsilon_a\tilde{\mu}), \quad \dot{\tau} = \frac{1}{2} \tag{15a}$$

$$\dot{\mu} = \frac{1}{\varepsilon_p}\mathcal{R}_\kappa\mu, \quad \mu \in \mathbb{T}^n, \tag{15b}$$

and discrete-time dynamics given by:

$$x^+ = x, \quad p^+ = r_0 x + (1 - r_0)p, \quad \tau^+ = \delta, \ \mu^+ = \mu, \tag{16}$$

where $r_o \in \{0, 1\}$ is a parameter that defines the *reset policy* of the controller, and $\delta > 0$ is a tunable parameter. In (15), we used $\tilde{\mu} := [\mu_1, \mu_3, \ldots, \mu_{2n-1}]^\top$ to denote the odd components of the state μ, which, by direct computation are $\mu_i(t) = \mu_i(0)\cos\left(\frac{2\pi\varepsilon_o\kappa_i t}{\varepsilon_p}\right) + \mu_{i+1}(0)\sin\left(\frac{2\pi\varepsilon_o\kappa_i t}{\varepsilon_p}\right)$, with $\mu_i(0)^2 + \mu_{i+1}(0)^2 = 1$. The continuous-time dynamics (15) are implemented whenever the states satisfy

$$(x, p, \tau, \mu) \in C_H \times \mathbb{T}^n, \quad C_H := \mathbb{R}^n \times \mathbb{R}^n \times [\delta, \Delta], \tag{17}$$

where $\Delta > \delta > 0$. On the other hand, the discrete-time dynamics (16) are implemented whenever the states satisfy

$$(x, p, \tau, \mu) \in D_H \times \mathbb{T}^n, \quad D_H := \mathbb{R}^n \times \mathbb{R}^n \times \{\Delta\}. \tag{18}$$

Therefore, it follows that the HAES implements periodic resets of the form (16) after intervals of flow of length $2(\Delta - \delta)$.

Lemma 1: [11, Lemma 7.2] Suppose that Assumption 4 holds. Then, the HAES has a well-defined average hybrid dynamical system with dynamics

$$\dot{x} = \frac{2}{\tau}(p - x), \ \dot{p} = -2k\tau\nabla\phi(x) + \mathcal{O}(\varepsilon_a), \ \dot{\tau} = \frac{1}{2}, \tag{19a}$$

$$x^+ = x, \quad p^+ = r_0 x + (1 - r_0)p, \quad \tau^+ = \delta, \tag{19b}$$

with flow set given by $\mathbb{R}^n \times \mathbb{R}^n \times [\delta, \Delta]$, and jump set given by $\mathbb{R}^n \times \mathbb{R}^n \times \{\Delta\}$. □

Remark 1: When $\delta = 0$ and $\Delta = \infty$, the average continuous-time dynamics (19) can be seen as perturbed version of (13) via the change of variables $\tau = t$, $p = x + \frac{1}{2}\tau\dot{x}$. Thus, it preserves the acceleration properties of (13) in each interval of flow. Moreover, note that the reset policy $r_0 = 1$ sets $p^+ = x$, which is equivalent to resetting the momentum $\dot{x}$ to zero. Similar resetting techniques, usually called *restarting mechanisms*, have been studied in the literature of optimization and machine learning; see [8]. □

The stability properties of the HAES (15)-(18) are studied with respect to the compact set

$$\mathcal{A}_s = \{(x, p) : x \in \mathcal{A}, \ p = x\} \times [\delta, \Delta] \times \mathbb{T}^n. \tag{20}$$

In particular, the following results were established in [11, Thm. 1-2] for the solution of model-free optimization problems of *static maps*.

Proposition 1: Suppose that Assumption 2 and 4 hold. Then, if $r_0 = 0$, the HAES renders the set $\mathcal{A}_s$ SGPAS as $(\varepsilon_p, \varepsilon_a) \to 0^+$. Moreover, for each compact set $K_0 \subset \mathbb{R}^{2n}$ such that $\mathcal{A} \subset \text{int}(K_0)$, and each $\nu > 0$, $\exists\ \varepsilon_a^* > 0$ such that $\forall\ \varepsilon_a \in (0, \varepsilon_a^*)$, $\exists\ \varepsilon_p^* > 0$ such that $\forall\ \varepsilon_p \in (0, \varepsilon_p^*)$, all solutions with $[x(0,0), p(0,0)] \in K_0$ induce the bound:

$$\phi(x(t,j)) - \phi^* \leq \frac{4V(z(\underline{t}_j, j))}{k(t - \underline{t}_j)^2} + \nu, \tag{21}$$

for all $(t, j) \in \text{dom}(z)$ such that $t > \underline{t}_j$, where $\underline{t}_j = \min\{t : (t, j) \in \text{dom}(z)\}$, and $V(\cdot)$ is a Lyapunov function for the average hybrid system (19) that satisfies $\limsup_{j\to\infty} V(z(\underline{t}_j, j)) \leq \nu$. □

Remark 2: Property (21) describes a "semi-acceleration" property. Namely, after each reset j, the error in the cost function is minimized at a rate of $\mathcal{O}_j(1/(t - \underline{t}_j)^2)$ in each interval of flow[3]. However, the numerator in the bound (21) changes every time there is a jump. □

Proposition 2: Suppose that Assumptions 2, 3 and 4 hold. Let the parameters (k, δ, Δ) satisfy the inequality:

$$\Delta^2 - \delta^2 \geq \frac{1}{2\kappa k}. \tag{22}$$

Then, the set $\mathcal{A}_s$ is SGPES as $(\varepsilon_p, \varepsilon_a) \to 0^+$, and for each compact set $K_0 \subset \mathbb{R}^{2n}$ such that $\mathcal{A} \subset \text{int}(K_0)$, and each $\nu > 0$, $\exists\ \varepsilon_a^* > 0$ such that $\forall\ \varepsilon_a \in (0, \varepsilon_a^*)$, $\exists\ \varepsilon_p^* > 0$ such that $\forall\ \varepsilon_p \in (0, \varepsilon_p^*)$, all solutions with $[x(0,0), p(0,0)] \in K_0$, $\tau(0,0) = \delta$, and $x(0,0) = p(0,0)$, satisfy the bound

$$\phi(x(t,j)) - \phi^* \leq \alpha_0\tilde{\gamma}^j\big(\phi(x(0,0)) - \phi^*\big) + \nu, \tag{23}$$

for all $(t, j) \in \text{dom}(z)$ where $\tilde{\gamma} := \frac{1}{k\Delta^2}\left(\frac{1}{2\delta} + k\delta^2\right)$ and $\alpha_0 := \Delta^2/\delta^2$. □

Remark 3: The bound (23) establishes an exponential decrease of the error in the cost function during jumps. The key observation is that the constant $\tilde{\gamma}$ can be adjusted by tuning the resetting frequency of the algorithm via the parameters (δ, Δ). In particular, when $\delta \approx 0$, the choice $\Delta = e\sqrt{\frac{1}{2k\kappa}}$, guarantees convergence of order $\mathcal{O}(1/\sqrt{k\kappa}\log(1/\epsilon))$, where $\epsilon > 0$ is the precision [11, Sec. 3.2.1]. In fact, it can be shown that this choice is actually optimal in periodic restarting mechanisms that reset the momentum state [19]. However, since the constant κ is unknown, the tuning of the reset parameters in the HAES must be carried out by trial and error. □

[3]Note that $\underline{t}_j = 0$ when $j = 0$, i.e., during the first interval of flow.

D. Main Results: HAES for Dynamic Plants

We now present the main theoretical result of this paper, which establishes the stability and convergence properties of the interconnection between the HAES and the dynamic plant (2). Additionally, we incorporate linear dynamic filters with state ξ and matrices $(A, B, \tilde{C})$, which have been shown to be instrumental for the implementation of ESC in practical applications [3]. The proof is omitted due to space limitations.

Theorem 1: Suppose that Assumptions 1,2, and 4 hold. Consider the closed-loop HDS dynamics with state $\vartheta = (\theta, \xi, z)$, $z = (x, p, \tau, \mu)$, and dynamics

$$\varepsilon_o \frac{d\theta}{ds} = f(\theta, x + \varepsilon_a \tilde{\mu}), \tag{24a}$$

$$\varepsilon_f \frac{d\xi}{ds} = A\xi + \frac{2}{\varepsilon_a} Bh(\theta, x + \varepsilon_a \tilde{\mu})\tilde{\mu}, \tag{24b}$$

$$\varepsilon_p \frac{d\mu}{ds} = \mathcal{R}_\kappa \mu, \quad \mu \in \mathbb{T}^n, \tag{24c}$$

$$\frac{dx}{ds} = \frac{2}{\tau}(p - x), \quad \frac{dp}{ds} = -2k\tau \tilde{C}\xi, \quad \frac{d\tau}{ds} = \frac{1}{2}, \tag{24d}$$

evolving in the flow set

$$C := \mathbb{R}^p \times \mathbb{T}^n \times \mathbb{R}^n \times C_H, \tag{25}$$

and discrete-time dynamics:

$$\theta^+ = \theta, \quad \mu^+ = \mu, \quad \xi^+ = \xi, \quad x^+ = x, \tag{26a}$$

$$p^+ = r_0 x + (1 - r_0)p, \quad \tau^+ = \delta, \tag{26b}$$

evolving in the jump set

$$D := \mathbb{R}^p \times \mathbb{T}^n \times \mathbb{R}^n \times D_H. \tag{27}$$

Then, if A is Hurwitz, and $-\tilde{C}A^{-1}B = I$, we have:

1) If $r_o = 0$, the set $\ell(\mathcal{A}_s) \times \{0\} \times \mathcal{A}_s$ is SGPAS as $(\varepsilon_o, \varepsilon_p, \varepsilon_a, \varepsilon_f) \to 0^+$, and the bound (21) holds in the (s, j) time scale.
2) If $r_o = 1$, Assumption 3 holds, and condition (22) is satisfied, then the set $\ell(\mathcal{A}_s) \times \{0\} \times \mathcal{A}_s$ is SGPAS as $(\varepsilon_o, \varepsilon_p, \varepsilon_a, \varepsilon_f) \to 0^+$, and the bound (23) holds in the (s, j) time scale. □

Remark 4: Theorem 1 is analogous to the SGPAS results established in the literature for gradient descent-based ESC [20, Thm. 1], for Newton-like ESC [15, Thm. 1], and for Fixed-Time ESC [16, Thm. 1]. □

Remark 5: Note that the bounds (21)- (23) only hold in the limit as $(\varepsilon_o, \varepsilon_p, \varepsilon_a, \varepsilon_f) \to 0^+$, and in the (s, j) time scale. Therefore, the transient performance of the ESC is naturally limited by the transient performance of the plant. □

IV. Traffic Light Control via Hybrid Accelerated Extremum Seeking

In this section, we apply the HAES to optimize self-organizing traffic lights (SOLT) systems for *signalized urban intersections*, a class of systems explained in detail in [12]. At each intersection, the possible traffic flows are sorted into *phases*. Then, the flows in the same phase receive a green light simultaneously and are not interrupted until they receive a red light. All the phases at the intersection are hence coordinated in a *cycle*. SOTL receives measurements in real-time and sets a control decision based on the value of a *threshold function* u of each phase of the intersection. When the value of this function exceeds a threshold value, the phase becomes a candidate to be the next active phase. This threshold function quantifies how congested the phase is, and therefore it affects the performance of the controller. Fig. 1 illustrates the impact of the threshold on the performance of the network. Here, the parameter-average speed map exhibits "almost flat" regions where the convergence of standard gradient descent-based algorithms can be prohibitively slow. To address this issue, we implement the HAES, as shown in Fig. 2, which corresponds to (15) with $\mu(0, 0) = [0, 1]$, $a = \varepsilon_a$ and $\omega = \frac{2\pi\kappa}{\varepsilon_p \varepsilon_0}$. The disturbance models small measurement noise, unavoidable in practical applications.

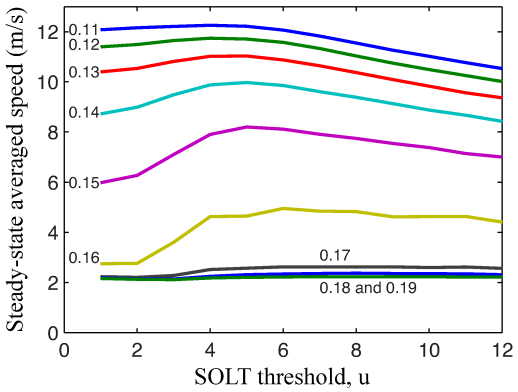

Fig. 1: Response maps for different demands, ranging from light traffic ($t_r = 0.11$) to heavy congestion ($t_r = 0.19$) [2].

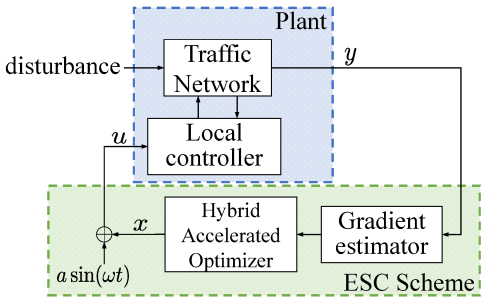

Fig. 2: Scheme of Hybrid Accelerated Extremum Seeking Control applied to a traffic network.

To simulate the traffic network we used a cellular automata model as in [12]. The traffic network that we study is a square grid with 16 intersections, as shown in Fig. 3. The simulation setup of the network is the same as the one used in [2]. At each time step, a vehicle enters the network with probability α_i and leaves it with probability β_i, thus simulating perimeter or boundary conditions of the local network. Furthermore, there are four phases at each intersection: an east/west phase, an east/west turning phase, a north/south phase, and a north/south turning phase. Each

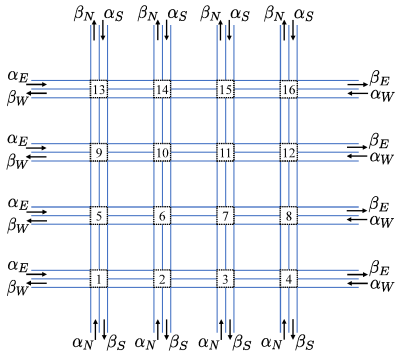

Fig. 3: The traffic network considered in the simulation.

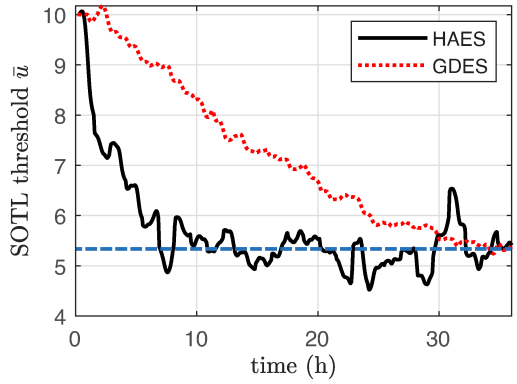

Fig. 4: Comparison between convergence of HAES and GDES in SOLT.

vehicle makes a turning decision based on a predetermined probability. The state of this traffic network model is the queue length on the road, the dynamics of this state are modeled by the conservation of vehicles, and flow equations which consider a nonlinear flow rate that captures different events, like intersection queuing, lane changing, turning, among others [2]. Moreover, the traffic flows from/to outside the network is considered as disturbances. The traffic light setting is controlled by SOTL. The ESC receives the network average speed y as an input, and generates the threshold value u as an output. Fig. 4 shows the simulation results with $\alpha_i = 0.15$ and $\beta_i = 0.9$ (which corresponds to the purple line in Fig. 1 indicated by the label 0.15). It shows the convergence of x by using both HAES and the gradient descent-based ES (GDES). The comparison was done by using equivalent parameters for both schemes. The common parameters are: $\varepsilon_a = 1$, $\omega = 10^{-3}$ rad/s, a gain of 0.35, and a band-pass filter with cut-off frequencies at 9×10^{-5} rad/s and 3×10^{-3} rad/s. Additionally for the HAES, $\delta = 0.1$ and $\Delta = 2.5$. We also added a phase lag of 60° into the controller to compensate for the dynamics of the plant. This phase lag was shown to reduce the steady state oscillations. As it can be observed, given an equivalent set of parameters, the HAES is able to converge to a neighborhood of the optimal point significantly faster compared to the traditional GDES.

V. Conclusions

We presented two main results. First, we showed that the hybrid accelerated extremum seeking (HAES) dynamics interconnected with a stable dynamic plant in the loop preserves the semi-global practical asymptotic stability properties, and also preserves the acceleration bounds as the time-scale separation increases. Second, we implemented the HAES to optimize a self-organizing traffic system (SOTL), and we showed via simulation that the HAES can substantially decrease the convergence time. This application provides a suitable example of the importance of dynamic momentum in ESC to improve transient performance in real-time model-free optimization problems.

References

[1] H. Wei, G. Zheng, V. Gayah, and Z. Li, "A survey on traffic signal control methods," *arXiv:1904.08117*, 2019.

[2] R. Kutadinata, W. Moase, C. Manzie, L. Zhang, and T. Garoni, "Enhancing the performance of existing urban traffic light control through extremum-seeking," *Transportation Research Part C: Emerging Technologies*, vol. 62, pp. 1–20, 2016.

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

[4] D. Nešić, Y. Tan, W. H. Moase, and C. Manzie, "A unifying approach to extremum seeking: Adaptive schemes based on estimation of derivatives," *49th IEEE Conference on Decision and Control*, pp. 4625–4630, 2010.

[5] J. I. Poveda and A. R. Teel, "A framework for a class of hybrid extremum seeking controllers with dynamic inclusions," *Automatica*, vol. 76, pp. 113–126, 2017.

[6] J. I. Poveda, R. Kutadinata, C. Manzie, D. Nešíc, A. R. Teel, and C. Liao, "Hybrid extremum seeking for black-box optimization in hybrid plants: An analytical framework," *57th IEEE Conferece on Decision and Control*, pp. 2235–2240, 2018.

[7] A. C. Wilson, B. Recht, and M. I. Jordan, "A Lyapunov analysis of momentum methods in optimization," *Journal of Machine Learning Research*, no. 22, pp. 1–34, 2021.

[8] W. Su, S. Boyd, and E. Candes, "A differential equation for modeling Nesterov's accelerated gradient method: Theory and insights," *Journal of Machine Learning Research*, vol. 17, no. 153, pp. 1–43, 2016.

[9] A. Wibisono, A. C. Wilson, and M. I. Jordan, "A variational perspective on accelerated methods in optimization," *Proceedings of the National Academy of Sciences*, vol. 113, no. 47, pp. E7351–E7358, 2016.

[10] J. I. Poveda and N. Li, "Inducing uniform asymptotic stability in time-varying accelerated optimization dynamics via hybrid regularization," *58th IEEE Conference on Decision and Control*, pp. 3000–3005, 2019.

[11] ——, "Robust hybrid zero-order optimization algorithms with acceleration via averaging in time," *Automatica*, vol. 123, p. 109361, 2021.

[12] J. de Gier, T. M. Garoni, and O. Rojas, "Traffic flow on realistic road networks with adaptive traffic lights," *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2011, no. 04, p. P04008, apr 2011.

[13] R. Goebel, R. Sanfelice, and A. R. Teel, *Hybrid Dynamical System*. Princeton, NJ: Princeton University Press, 2012.

[14] Y. Tan, D. Nešić, and I. Mareels, "On non-local stability properties of extremum seeking controllers," *Automatica*, vol. 42, no. 6, pp. 889–903, 2006.

[15] A. Ghaffari, M. Krstić, and D. Nešić, "Multivariable newton-based extremum seeking," *Automatica*, vol. 48, pp. 1759–1767, 2012.

[16] J. I. Poveda and M. Krstic, "Non-smooth extremum seeking control with user-prescribed fixed-time convergence," *IEEE Transactions on Automatic Control*, pp. 1–1, 2021.

[17] Y. Nesterov, *Introductory Lectures on Convex Optimization: A Basic course*. Boston, MA.: Kluwer Academic Publishers, 2004.

[18] J. I. Poveda and A. R. Teel, "The heavy-ball ode with time-varying damping: Persistence of excitation and uniform asymptotic stability," *In Proc. of American Control Conference.*, pp. 773–778, 2020.

[19] O'Donoghue and E. J. Candes, "Adaptive restart for accelerated gradient schemes," *Foundations of Computational Mathematics*, vol. 15, no. 3, pp. 715–732, 2013.

[20] M. Krstić and H.-H. Wang, "Stability of extremum seeking feedback for general nonlinear dynamic systems," *Automatica*, vol. 36, no. 4, pp. 595–601, 2000.