The Internal Model Principle of Time-Varying Optimization

Gianluca Bianchin
Bryan Van Scoy

***Abstract*—Time-varying optimization problems are central to many engineering applications where performance metrics and system constraints evolve dynamically with time. A number of algorithms have been proposed in recent years to solve such problems; a common feature of all these methods is that they implicitly require precise knowledge of the temporal variability of the solutions in order to exactly track the optimizers. In this paper, we seek to lift these stringent assumptions. Our main result is a fundamental characterization showing that an algorithm can track an optimal trajectory if and only if it contains a model of the temporal variability of the problem. We refer to this concept to as the *internal model principle of time-varying optimization.* By recasting the optimization objective as a nonlinear regulation problem and using tools from center manifold theory, we provide necessary and sufficient conditions both for an optimization algorithm to achieve exact asymptotic tracking and for such an algorithm to exist. We illustrate the applicability of the approach numerically on both synthetic problems as well as practical problems in transportation.**

I. Introduction

Time-varying optimization problems play a central role in several scientific domains, as they underpin many important contemporary engineering problems. Examples include training in Machine Learning [1], [2], dynamic signal estimation in Signal Processing [3], trajectory tracking in Robotics [4], system optimization in Industrial Control [5], and much more. Historically, discrete-time algorithms for time-varying optimization have been proposed and studied first since they emerge as a natural extension of their time-invariant counterparts, allowing for cost functions and constraints that may change over time [6]. These approaches build on the classical perspective on mathematical optimization, which seeks to construct methods to determine optimizers and consist of iterative procedures implemented on digital devices. Recently, there has been a renewed interest in the use and development of continuous-time dynamics for optimization purposes, due to the possibility of utilizing tools from dynamical systems for their analysis [7], [8] and motivated by practice, where optimization is increasingly used to control physical systems [9].

Motivated by these recent developments, in this paper we study time-varying convex optimization problems and focus on the use of continuous-time dynamics to track exactly optimal solutions. Although, by now, several approaches have been developed for this purpose [10]–[13], all these techniques implicitly require full knowledge of the temporal variability of the problem [13]. Unfortunately, in most practical applications, having such knowledge is impractical, either because the temporal variability enters the optimization in the form of exogenous disturbances that are unknown and cannot be measured (see, e.g., [9], [14]), or simply because it is unrealistic to ask for a noiseless model of how the problem changes with time. Departing from this, in this paper we pose the following question: is it possible to track (exactly and asymptotically) a minimizer of a time-varying optimization problem without any knowledge of the temporal variability of the optimization? Interestingly, we prove a fundamental result showing that tracking can be achieved if and only if the temporal variability of the problem can be 'observed' by the algorithm, and the latter incorporates a suitably reduplicated model of such a variability. We refer to this conclusion as *the internal model principle of time-varying optimization,* akin to its control-theoretic counterpart [15], [16]. Our approach relies on reinterpreting the optimization algorithm design as a nonlinear, multivariable regulation problem [17], and our analysis uses tools from center manifold theory [16], [18]. Figure 1 illustrates the architecture of the gradient-feedback design scheme studied in this paper.

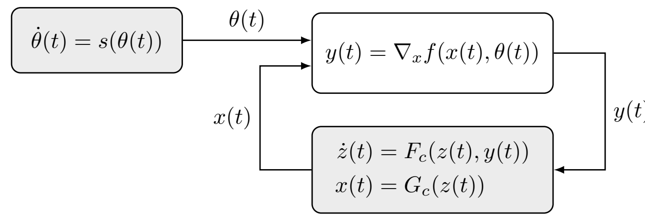

Fig. 1: Architecture of the gradient-feedback design scheme studied in this work. An optimization algorithm is to be designed (bottom block), having access only to gradient evaluations of the loss function to be minimized (top right block), and generating a sequence of exploration points $x(t)$ at which the gradient shall be evaluated. The loss function to be optimized varies with time, where the temporal variability $\theta(t)$ is assumed to be unmeasurable and generated by an exosystem (top left block). Shaded blocks emphasize the presence of dynamics.

Related works. The literature on methods for time-varying optimization is mainly divided into two classes of solutions. The first class consists of methods that do not utilize any model of the temporal variability of the problem [19]–[22]; instead, they seek to solve a sequence of static problems. Several established approaches belong to this class, including the online gradient descent method [23] and the online Newton step algorithm [24] – see [25] and references therein. Because the temporal variability of the problem is unknown (or ignored), algorithms in this class can make a decision only *after* each variation has been observed, thus incurring a certain 'regret.' Mathematically, these techniques are capable of reaching only a neighborhood of an optimizer [13], and exact tracking is out of reach in general for these approaches.

G. Bianchin is with the ICTEAM institute and the Department of Mathematical Engineering (INMA) at the University of Louvain, Belgium. Email: gianluca.bianchin@uclouvain.be

B. Van Scoy is with the Dept. of Electrical and Computer Engineering, Miami University, OH 45056, USA. Email: bvanscoy@miamioh.edu

In contrast, the second class of methods uses a model of the temporal evolution of the problem to seek to exactly track optimal trajectories [2], [26]–[28]. Particularly celebrated is the prediction-correction algorithm (see [10] and the recent work [11]), whereby at each time a prediction step is used to anticipate how the optimizer will evolve over time, and a correction step is used to seek a solution to each instantaneous optimization problem. We refer to [13] for a recent overview of the topic. Recent years have witnessed a growing interest in this problem: [4] uses contraction to study these methods; [29] uses sampling to estimate the temporal variability of the problem; a recent survey on analyzing optimization algorithms using tools from control has appeared in [30]; constrained optimization problems are studied in [31]. Despite these efforts, the fundamental question of what is the basic algorithm structure that enables exact asymptotic tracking remains unanswered.

Contributions. This paper makes four main contributions. First, we recast the problem of designing a time-varying optimization algorithm as a nonlinear multivariable regulation problem, whereby an algorithm is to be designed to regulate the gradient of a certain function to zero. We leverage this formulation to characterize the class of optimization algorithms that can achieve tracking. In a net departure from existing approaches (e.g., [11]–[13], [23]), our characterization is general and allows us to study not only a single optimization method, but an entire class, which enables us to derive fundamental results for all algorithms in the class. Second, by harnessing tools from center manifold theory [16], [18], we provide necessary and sufficient conditions for an optimization algorithm to ensure tracking. Interestingly, these conditions depend on the properties of the loss function (through a gradient invertibility-type condition) and on the inner model describing the temporal variability of the problem. This property allows us to prove the *internal model principle of time-varying optimization,* which states that for an optimization algorithm to achieve asymptotic tracking, it must incorporate a reduplicated model of the temporal variability of the problem. This feature is implicit in all existing approaches for time-varying optimization [13] but, to the best of the authors' knowledge, lacked a formal proof until now. Third, we derive necessary and sufficient conditions for the existence of a tracking algorithm. Fourth, we use our characterizations to design algorithms for time-varying optimization. With respect to the existing literature, our algorithm does not require one to know or measure exactly the temporal variability of the problem, and thus uses less stringent assumptions. Finally, we illustrate the applicability of the approach numerically on both synthetic problems as well as practical problems in transportation.

Organization. Section II presents the problem of interest, Section III studies the simpler problem of parameter-feedback optimization, which allows us to derive the necessary framework, and Section IV tackles the time-varying optimization design problem in its full generality and contains our main results. Section V discusses the tracking accuracy in relationship to the fidelity of the internal model, Section VI presents some extensions to constrained optimization, Section VII validates numerically the results, and Section VIII illustrates our conclusions. Finally, in the Appendix we summarize basic facts on center manifold theory that are used extensively throughout the paper.

Notation. We denote by $\mathbb{S}^n$ the space of $n \times n$ symmetric real matrices. Given an open set U, we say that $f : U \to \mathbb{R}$ is of differentiability class C^k if it has a k^{th} derivative that is continuous in U. The gradient of $f(x,\theta) : \mathbb{R}^n \times \Theta \to \mathbb{R}, \Theta \subseteq \mathbb{R}^p$, with respect to $x \in \mathbb{R}^n$ is denoted by $\nabla_x f(x,\theta) : \mathbb{R}^n \times \Theta \to \mathbb{R}^n$. The partial derivatives of $\nabla_x f(x,\theta)$ with respect to x and θ are denoted by $\nabla_{xx} f(x,\theta) : \mathbb{R}^n \times \Theta \to \mathbb{S}^n$ and $\nabla_{x\theta} f(x,\theta) : \mathbb{R}^n \times \Theta \to \mathbb{R}^{n\times p}$, respectively.

II. Problem setting

A. Problem statement

We consider the time-varying optimization problem:

$$\min_{x\in\mathbb{R}^n} f(x,\theta(t)), \tag{1}$$

where $t \in \mathbb{R}_{\geq 0}$ denotes time and $f : \mathbb{R}^n \times \Theta \to \mathbb{R}$, $\Theta \subseteq \mathbb{R}^p$, is a loss function that is parametrized by the time-varying parameter vector $\theta : \mathbb{R}_{\geq 0} \to \Theta$. We make the following assumptions throughout.

Assumption 1 (Properties of the objective function)**.** The map $x \mapsto f(x,\theta)$ is convex and $x \mapsto \nabla_x f(x,\theta)$ is Lipschitz continuous in $\mathbb{R}^n$, for each $\theta \in \Theta$. □

Assumption 2 (Existence of an exosystem)**.** There exists a smooth (i.e., C^∞) vector field $s : \Theta \to \mathbb{R}^p$ and $\theta(0) \in \Theta$ such that the parameter vector $\theta(t)$ satisfies

$$\dot\theta(t) = s(\theta(t)) \tag{2}$$

for all $t \in \mathbb{R}_{\geq 0}$. □

Convexity and smoothness are standard assumptions in optimization [25]. Assumption 2 ensures the existence of an autonomous system, called *exosystem*, describing the class of temporal variabilities of the cost taken into consideration.

In what follows, we will say that $x^\star : \mathbb{R}_{\geq 0} \to \mathbb{R}^n$ is a *critical trajectory* of (1) if it satisfies:

$$0 = \nabla_x f(x^\star(t), \theta(t)), \quad \forall t \in \mathbb{R}_{\geq 0}.$$

We will assume existence of a critical trajectory and that any critical trajectory is continuous.

Remark 1. The time variation in (1) is captured implicitly through the parameter vector $\theta(t)$. A related, yet slightly more general, problem is as follows:

$$\min_{x\in\mathbb{R}^n} f_0(x,t), \tag{3}$$

where the dependency on time is explicit. Problem (1) can be recast uniquely as in (3) by letting $f_0(x,t) = f(x,\theta(t))$ for all x and t. On the other hand, in general, there exists an infinite number of ways to parametrize (3) as in (1), thus leading to possible ambiguities. For instance, any function $f_0(x,t)$ may be parametrized by $\theta(t) = t$ (so that $f_0 \equiv f$), although this is not always compatible with our assumptions (see Assumption 3 and Remark 2 below). □

In this work, we are interested in the problem of designing an optimization algorithm that is capable of determining and tracking a critical trajectory of (1). We are interested in doing so with as little knowledge as possible on $\theta(t)$ and its temporal variability. From an optimization perspective, three types of assumptions may be considered (presented in increasing order of restrictiveness):

(A1) The parameter signal $\theta(t)$ is known or measurable at each $t \in \mathbb{R}_{\geq 0}$, but no knowledge of $s(\theta)$ is available.
(A2) The vector field $s(\theta)$ is known, but $\theta(0)$ (and hence also $\theta(t)$) is unknown.
(A3) The vector field $s(\theta)$ as well as $\theta(0)$ are known.

Under these assumptions, different methods have been developed in the literature, each with corresponding guarantees:

(M1) Under (A1), all methods from static optimization (e.g., [13], [19]–[24]) can be applied directly. All these methods can achieve convergence only to a neighborhood of a critical trajectory.
(M2) The approach of [12] is applicable in this case, which is however limited to strongly convex quadratic optimization problems with linear temporal variablities.
(M3) The class of prediction correction methods [10] and their variations [11] take advantage of this knowledge and ensures asymptotic tracking of $x^\star(t)$.

A natural question emerges directly from the above discussion: is (A3) necessary to ensure exact tracking and, if not, can it be relaxed (e.g., to (A1) or (A2))? With this motivation, in this paper we are driven by the following question.

Problem 0 (Minimal knowledge for exact tracking)**.** What is the least-restrictive set of assumptions (cf. (A1)–(A3)) under which there exists an optimization algorithm that achieves exact asymptotic tracking of the critical points of (1)? □

Our objective is to seek an optimization algorithm that assumes no access to $\theta(t)$; instead, as common in first-order optimization approaches [25], we will assume only the availability of functional evaluations of the gradient function $\nabla_x f(x, \theta(t))$ at points $x \in \mathbb{R}^n$, selected by the algorithm. Formally, the optimization algorithm will be described by an internal state $z(t)$, which takes values on an open subset $\mathcal{Z} \subseteq \mathbb{R}^{n_c}, n_c \in \mathbb{N}_{>0}$; the optimization algorithm generates a sequence of points $x(t) \in \mathbb{R}^n$ (called *exploration signal*) at which the gradient shall be evaluated, and processes functional evaluations of the gradient at these points, $y(t) = \nabla_x f(x(t), \theta(t))$ (called *gradient feedback signal*). Mathematically, the optimization algorithm is:

$$\dot{z}(t) = F_c(z(t), y(t)), \qquad x(t) = G_c(z(t)), \tag{4a}$$

together with the *gradient-feedback signal*:

$$y(t) = \nabla_x f(x(t), \theta(t)), \tag{4b}$$

where $F_c : \mathcal{Z} \times \mathbb{R}^n \to \mathbb{R}^{n_c}$ and $G_c : \mathcal{Z} \to \mathbb{R}^n$ are functions to be designed. In the remainder, we refer to (4) as a *dynamic gradient-feedback optimization algorithm*. The architecture of the studied gradient-feedback algorithm is illustrated in Figure 1. Notice that the dynamics of the optimization algorithm (4), coupled with the time-variability generator (2), have the form of a nonlinear autonomous system:

$$\begin{aligned}\dot{z}(t) &= F_c(z(t), y(t)),\\ y(t) &= \nabla_x f(G_c(z(t)), \theta(t)),\\ \dot{\theta}(t) &= s(\theta(t)).\end{aligned} \tag{5}$$

Our objective is to design $F_c(z, y)$, $G_c(z)$, and n_c, so that $y(t) \to 0$ as $t \to \infty$, which ensures that $x(t)$ tracks, with zero asymptotic error, a critical trajectory $x^\star(t)$ of (1).

We will assume that $\theta(t)$ takes values in a neighborhood of the origin, and thus let Θ be some neighborhood of the origin of $\mathbb{R}^p$. Note that there is no loss of generality in doing so, because if $\theta(t)$ takes values in the neighborhood of any other point, the former can be shifted to the origin via a change of variables without altering the critical points of (1). In what follows, we will denote by $x^\star_\circ \in \mathbb{R}^n$ a point such that

$$0 = \nabla_x f(x^\star_\circ, 0), \tag{6}$$

and assume that $x^\star_\circ \in \mathbb{R}^n$ is locally unique. Moreover, we will let $z^*_\circ$ be such that $x^\star_\circ = G_c(z^\star_\circ)$, and assume that $z^*_\circ$ exists and is locally unique.

Definition 1. We say that (4) *asymptotically tracks a critical trajectory* of (1) if, for each initial condition $(z(0), \theta(0))$ in some neighborhood of $(z^*_\circ, 0) \in \mathcal{Z} \times \Theta$, the solution of (5) satisfies $y(t) \to 0$ as $t \to \infty$. □

Notice that, when an algorithm asymptotically tracks a critical trajectory, $x(t) \to x^\star(t)$ for some critical trajectory $x^\star(t)$; namely, critical point is readily given by the exploration signal. We make the objective of this work formal as follows.

Problem 1 (Dynamic gradient-feedback problem)**.** Find necessary and sufficient conditions (in terms of the loss function f) for the existence of a gradient-feedback optimization algorithm that asymptotically tracks a critical trajectory of (1). Derive a method to design such an algorithm when these conditions hold. □

It is important to note that existence conditions for $F_c(z, y)$ and $G_c(z)$ shall not depend on the temporal variability of the optimization (namely, $s(\theta)$), but only on the properties of the optimization problem. Conversely, if such conditions were to depend on $s(\theta)$, a more general class of optimization algorithms than (4) could be constructed, implying the formulation (4) would not be sufficiently general.

B. Standing assumptions

We list hereafter the basic assumptions on which our approach to gradient-feedback theory is based.

Assumption 3 (Stability of the exosystem)**.** The point $\theta = 0$ is a stable equilibrium of (2), and there exists a neighborhood $\Theta_\circ \subset \Theta$ of 0 with the property that each initial condition $\theta(0) \in \Theta_\circ$ is Poisson stable[1]. □

Intuitively, $\theta(0)$ is Poisson stable if the corresponding trajectory returns to an arbitrarily small neighborhood of each of its points an infinite number of times. More explicitly, Assumption 3 implies that the matrix $S := \left[\frac{\partial s}{\partial \theta}\right]_{\theta=0}$, which characterizes the linear approximation of the exosystem, has all eigenvalues on the imaginary axis. In fact, no eigenvalue of S can have positive real part because otherwise the equilibrium would be unstable. Moreover, no eigenvalue can have negative real part, otherwise the exosystem would admit a stable invariant manifold near the equilibrium, and trajectories originating on this manifold would converge to zero as time tends to infinity, thus violating Poisson stability. The class of exosystems satisfying Assumption 3 includes the (important in practice) class of systems in which every solution is periodic. Moreover, any stable modes of $\theta(t)$ that converge to the origin asymptotically may be disregarded without changing the asymptotic behavior of the algorithm. Assumption 3 is instrumental for proving that certain conditions are necessary for the existence of a gradient-feedback algorithm, and can be dispensed if one is interested only in sufficient conditions; moreover, the assumption may be relaxed when the exosystem is linear (see Remark 2 below).

For convenience, we will require that $F_c(z_\circ^\star, 0) = 0$. This ensures that the optimization algorithm (4) has an equilibrium at $z = z_\circ^\star$, and the corresponding gradient feedback signal is identically zero: $y(t) = \nabla_x f(x_\circ^\star, 0) = 0$. Since $\theta(t)$ is unmeasurable, the temporal variability of the cost can only be evaluated through measurements of $y(t)$; hence, we make the following assumption.

Assumption 4 (Detectability of the exosystem)**.** Let

$$Q := \left[\frac{\partial \nabla_x f}{\partial \theta}\right]_{(x,\theta)=(x_\circ^\star,0)}, \qquad S := \left[\frac{\partial s}{\partial \theta}\right]_{\theta=0}. \quad (7)$$

The pair (Q, S) is detectable. □

Detectability ensures that all the modes of the exosystem (2) can be reconstructed from the available measurements $y(t)$. Undetectability of (Q, S) corresponds to a redundant description of the exogenous signal: indeed, if some modes of the exosystem do not influence the gradient, then $y(t)$ would be independent of those modes and they can thus be removed without altering the problem.

[1]Recall that, for a nonlinear system of the form (2), an initial condition θ° is said to be Poisson stable if the solution starting from θ° at time t of (2) (here denoted by $\theta_t(\theta^\circ))$) is defined for all $t \in \mathbb{R}_{\geq 0}$ and, for every neighborhood U of θ° and each $T \in \mathbb{R}_{>0}$, there exists a time $t_1 > T$ such that $\theta_{t_1}(\theta^\circ) \in U$, and a time $t_2 < -T$ such that $\theta_{t_2}(\theta^\circ) \in U$.

Example 1 (Detectability)**.** Suppose the objective function admits the decomposition $f(x, \theta(t)) = \hat{f}(x) + x^\mathsf{T}\theta_1(t)$, where $\theta_1(t) = \alpha\cos(\omega t)$ for some α and ω. The parameter signal can be generated by (2) with initial condition $\theta(0) = (\alpha, 0)$ and vector field:

$$s(\theta) = \begin{bmatrix} 0 & 1 \\ -\omega^2 & 0 \end{bmatrix} \theta.$$

Notice that, although the gradient signal $y(t) = \nabla_x \hat{f}(x) + \theta_1(t)$ is independent of $\theta_2(t)$, both modes of the exosystem affect the gradient (indeed, $\theta_1(t) = \alpha\cos(\omega t)$ cannot be generated by a linear system with a single mode). Notice also that the pair (Q, S) is detectable. □

III. The parameter-feedback problem

To address Problem 1, we first study a simpler problem that allows us to derive the necessary framework to tackle our objectives in their generality. To this end, in place of the dynamic optimization algorithm (4), we begin by considering an algebraic optimization algorithm of the form:

$$x(t) = H_c(\theta(t)), \quad (8)$$

where $H_c : \Theta \to \mathbb{R}^n$ is a mapping to be designed; we will require that H_c is of class C^0 and satisfies the fixed-point condition $x_\circ^\star = H_c(0)$. Because of the explicit dependence on $\theta(t)$, we will refer to (8) to as a *parameter-feedback* optimization algorithm. Our objective is to design the map H_c so that the composition of (2), (4b), and (8):

$$\begin{aligned} y(t) &= \nabla_x f(H_c(\theta(t)), \theta(t)), \\ \dot{\theta}(t) &= s(\theta(t)), \end{aligned} \quad (9)$$

tracks, with zero asymptotic error, the critical trajectories of (1). For the framework considered in this section, Problem 1 is reformulated as follows.

Problem 2 (Static parameter-feedback problem)**.** Find necessary and sufficient conditions (in terms of the loss function f) for the existence of a parameter-feedback optimization algorithm that asymptotically tracks a critical trajectory of (1). Derive a method to design such an algorithm when these conditions hold. □

Solvability of the static parameter-feedback problem will depend on the existence of a function that zeros the gradient.

Definition 2 (Mapping zeroing the gradient)**.** We say that a mapping $H_c : \Theta \to \mathbb{R}^n$ *zeros the gradient* at the point $\theta \in \Theta$ if

$$0 = \nabla_x f(H_c(\theta), \theta). \quad (10)$$

Moreover, we say that H_c zeros the gradient *globally* if (10) holds for all $\theta \in \Theta$, and *locally* if (10) holds for all $\theta \in \Theta_\circ$, where $\Theta_\circ \subseteq \mathbb{R}^p$ is some neighborhood of the origin[2]. □

The following result characterizes all parameter-feedback optimization algorithms that achieve asymptotic tracking of a critical trajectory.

[2]The neighborhoods in Defn. 2 and Assumption 3 need not be the same.

Theorem 1 (Parameter-feedback algorithm characterization)**.** Let Assumptions 1, 2, and 3 hold. The parameter-feedback algorithm (9) asymptotically tracks a critical trajectory of (1) if and only if the mapping H_c zeros the gradient locally.

Proof. (Only if) Suppose $y(t) \to 0$ as $t \to \infty$. We will show that H_c zeros the gradient locally. By Poisson stability of the exosystem (cf. Assumption 3), there exists a neighborhood $\Theta_\circ \subset \Theta$ of the origin such that, for every $\theta_\circ \in \Theta_\circ$, every $\varepsilon > 0$, and every $T > 0$, the trajectory of (9) satisfies

$$\|\theta(t) - \theta_\circ\| < \varepsilon,$$

at some $t > T$. Then, $\lim_{t\to\infty} y(t) = 0$ can hold only if (10) holds for all $\theta \in \Theta_\circ$.

(If) Suppose H_c zeros the gradient on some neighborhood $\Theta_0 \subset \Theta$ of the origin, i.e., (10) holds for all $\theta \in \Theta_0$. By Assumption 3, the point $\theta = 0$ is a stable equilibrium of the exosystem, so there exists some other neighborhood $\Theta_1 \subset \Theta$ of the origin and time $T > 0$ such that $\theta(t) \in \Theta_0$ for all $t > T$ when $\theta(0) \in \Theta_1$. Since H_c zeros the gradient on Θ_0, this implies that $y(t) \to \infty$ as $t \to \infty$. □

Remark 2. Poisson stability of the exosystem is needed for the gradient zeroing condition to be *necessary* for asymptotic tracking of a critical trajectory (but is not needed for sufficiency). This is because Poisson stability implies that all points sufficiently close to the origin are revisited infinitely often, so H_c must zero the gradient on a neighborhood of the origin. More generally, H_c must zero the gradient on all limit points of $\theta(t)$. □

The result provides a full characterization of all parameter-feedback algorithms that ensure exact asymptotic tracking. In words, $x = H_c(\theta)$ is a parameter-feedback algorithm if and only if it zeros the gradient everywhere in a neighborhood of the critical point. As a byproduct, the result also provides a necessary and sufficient condition for the solvability of the parameter-feedback problem: the problem is solvable if and only if the set of solutions to the system of equations $0 = \nabla_x f(x, \theta)$ can be expressed, everywhere in a neighborhood of the origin of Θ, as the graph of a function $x = H_c(\theta)$. Finally, the theorem provides an explicit form for the desired parameter-feedback algorithm (8): this is given by $x(t) = H_c(\theta(t))$; this algorithm ensures that, for all initial states $\theta(0)$ sufficiently close to the origin, $y(t) \to 0$ as $t \to \infty$. In this sense, Theorem 1 provides a complete answer to Problem 2. We illustrate the applicability of the result and the necessity of the provided condition in the following example.

Example 2. Consider an instance of (1) with $n = p = 1$,

$$f(x, \theta) = (x-1)^2(x+1)^2 + \frac{8}{3\sqrt{3}}x + \theta x. \tag{11}$$

See Fig. 2(a) for an illustration of this function. For $\theta = 0$, the optimization problem associated with (11) admits two critical points: $x^\star_{\circ,1} = \frac{1}{\sqrt{3}}$ and $x^\star_{\circ,2} = -\frac{2}{\sqrt{3}}$; indeed, it follows from

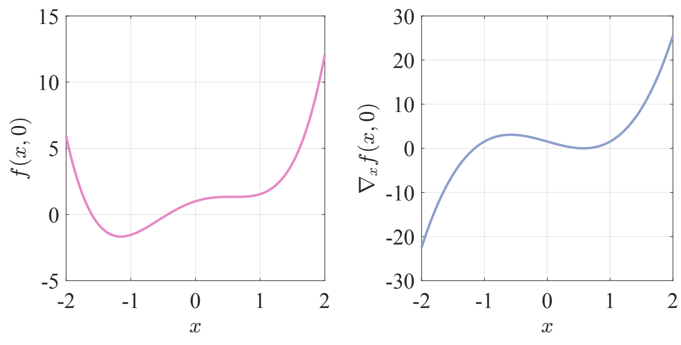

Fig. 2: Investigation of the condition (10). (Left) Loss function $f(x, \theta)$ studied in Example 2, plotted for $\theta = 0$. (Right) Gradient $\nabla f(x, \theta)$. The function $f(x, 0)$ admits two critical points: $x^\star_{\circ,1} = \frac{1}{\sqrt{3}}$ and $x^\star_{\circ,2} = -\frac{2}{\sqrt{3}}$. At $x^\star_{\circ,1}$, condition (10) is not satisfied, since with an an upward shift of the graph of $\nabla_x f(x, 0)$, $x^\star_{\circ,1}$ is no longer a critical point of $f(x, 0)$. On the other hand, (10) holds for $x^\star_{\circ,2}$, since $x^\star_{\circ,2}$ varies continuously as θ is perturbed. See Example 2 for a discussion.

direct inspection that $\nabla_x f(x^\star_{\circ,1}, 0) = 0$ and $\nabla_x f(x^\star_{\circ,2}, 0) = 0$. At the critical point $x^\star_{\circ,1}$, a function $H_c(\theta)$ as in (10) does not exist. This can be visualized with the aid of Fig. 2(b): if $\theta = 0$ is perturbed to $\theta+\epsilon$, $\epsilon > 0$, the graph of $\nabla_x f(x, 0)$ (illustrated in Fig. 2(b)) shifts upward and the equation $\nabla_x f(x, \epsilon) = 0$ no longer admits a solution in a neighborhood of $x^\star_{\circ,1}$.

On the other hand, for the critical point $x^\star_{\circ,2}$, any arbitrarily-small upward or downward shift of the graph of $\nabla_x f(x, 0)$ results in a continuous perturbation of $x^\star_{\circ,2}$ (see Fig. 2(b)), thus suggesting existence of $x = H_c(\theta)$ as in (10). This graphical observation can be formalized with the aid of the implicit function theorem [32], as described next. Define $F(x, \theta) := \nabla_x f(x, \theta)$ and notice that F is continuously differentiable with $F(x^\star_{\circ,2}, 0) = 0$. By the implicit function theorem, there exists a neighborhood $\Theta_\circ$ of $x^\star_{\circ,2}$ and a function $H_c : \Theta_\circ \to \mathbb{R}^n$ such that $F(H_c(\theta), \theta) = 0$ in $\Theta_\circ$, provided that $\frac{\partial F}{\partial x}\big|_{(x,\theta)=(x^\star_{\circ,2},0)} \neq 0$. By inspection, it is immediate to see that the latter condition is satisfied. □

Existence of a map $x = H_c(\theta)$ as in (10) can be ensured for general problems with the aid of the implicit function theorem [32], as illustrated next.

Proposition 2 (Existence of local parameter-feedback algorithms)**.** Let Assumptions 1, 2, and 3 hold, and let $X_\circ \times \Theta_\circ$ be some neighborhood of $(x^\star_\circ, 0)$. Further, assume that:

(i) the loss function f is C^1 on $X_\circ \times \Theta_\circ$,
(ii) $x \mapsto f(x, \theta)$ is C^2 on $X_\circ$ for each $\theta \in \Theta_\circ$, and
(iii) the Hessian $\nabla^2_{xx} f(x, \theta)\big|_{x=x^\star_\circ, \theta=0}$ is positive definite.

Then, there exits a C^0 mapping $H_c : \Theta \to \mathbb{R}^n$ that zeros the gradient locally. □

Proof. Define $F(x, \theta) = \nabla_x f(x, \theta)$ and note that, under the stated assumptions, F is C^0 on $X_\circ \times \Theta_\circ$, the mapping $x \mapsto$

$F(x,\theta)$ is C^1 on $X_\circ$ for all $\theta \in \Theta_\circ$, and

$$\det \left[\frac{\partial F(x,\theta)}{\partial x} \right]_{(x,\theta)=(x_\circ^\star,0)} \neq 0.$$

Then, the result follows from the implicit function theorem (see, e.g., [33, Thm. 1]) applied to the first-order optimality conditions $0 = F(x,\theta)$ at the point $(x,\theta) = (x_\circ^\star, 0)$. □

Intuitively, when the loss function is twice continuously differentiable in the decision variable with a positive definite Hessian, every critical point is a local isolated minimum and existence of $H_c(\theta(t))$ is guaranteed by the implicit function theorem. Notice that, differently from the characterization in Theorem 1, the conditions offered by Proposition 2 are sufficient but not necessary.

Although the conditions are immediate to verify, the convergence claims of Theorem 1 and Proposition 2 are of local nature, namely, $y(t) \to 0$ is ensured provided that $\theta(0)$ is sufficiently close to the origin. The following result provides a sufficient condition for global convergence.

Theorem 3 (Existence of global parameter-feedback algorithms)**.** Let Assumptions 1, 2, and 3 hold. Further, assume:

(i) the loss function f is C^1 on $\mathbb{R}^n \times \Theta$,
(ii) $x \mapsto f(x,\theta)$ is C^3 on $\mathbb{R}^n$ for each $\theta \in \Theta$,
(iii) the Hessian $\nabla_{xx}^2 f(x,\theta)$ is positive definite on $\mathbb{R}^n \times \Theta$,
(iv) the mappings $\nabla_{xx}^2 f$ and $\frac{\partial \nabla_{xx}^2 f}{\partial x}$ are C^0 on $\mathbb{R}^n \times \Theta$.

Then, there exits a C^0 mapping $H_c : \Theta \to \mathbb{R}^n$ that zeros the gradient globally. □

Proof. Define $F(x,\theta) = \nabla_x f(x,\theta)$ and note that, under the stated assumptions, F is C^0 on $\mathbb{R}^n \times \Theta$, the mapping $x \mapsto F(x,\theta)$ is C^2 on $\mathbb{R}^n$ for all $\theta \in \Theta$,

$$\det \left[\frac{\partial F(x,\theta)}{\partial x} \right] \neq 0, \quad \forall (x,\theta) \in \mathbb{R}^n \times \Theta,$$

and the mappings $\frac{\partial F}{\partial x}$ and $\frac{\partial^2 F}{\partial x^2}$ are C^0 in $\mathbb{R}^n \times \Theta$. Hence, assumptions (B1)–(B5) of a global version of the implicit function theorem [33, Thm. 2] (see also [34, Thm. 6]) are satisfied for the first-order optimality conditions $0 = F(x,\theta)$, and the claim follows. □

Theorem 3 shows that, under additional continuity assumptions on the loss function (cf. conditions (i), (ii), and (iv)) and when the Hessian of the loss is positive definite everywhere, convergence of (9) is ensured globally. More formally, under the assumptions of Theorem 3, the parameter-feedback algorithm $x(t) = H_c(\theta(t))$ ensures that $y(t) \to 0$ as $t \to \infty$ for all initial conditions $\theta(0) \in \Theta$. This result is particularly relevant, as it ensures the existence of parameter-feedback algorithms for the (very important in practice) class of strictly convex loss functions (studied in, e.g., [11], [12]).

We conclude this section by illustrating the design procedure for parameter-feedback algorithms on a quadratic problem.

Example 3. Consider an instance of (1) with quadratic cost and time-variability that depends linearly on $\theta(t)$ (which has been investigated in [12]):

$$f(x(t),\theta(t)) = \tfrac{1}{2} x(t)^\mathsf{T} R x(t) + x(t)^\mathsf{T} Q \theta(t), \qquad (12)$$

with matrices $R \in \mathbb{S}^n$ and $Q \in \mathbb{R}^{n\times p}$. In this case, the signal we wish to regulate to zero is: $y(t) = \nabla_x f(x(t),\theta(t)) = Rx(t) + Q\theta(t)$. For arbitrary θ, this problem admits a critical point if and only if $\operatorname{Im} Q \subseteq \operatorname{Im} R$, in which case $x_\circ^\star$ is unique. Applying Theorem 1 amounts to finding a linear transformation $H_c \in \mathbb{R}^{n\times p}$ such that $0 = (RH_c + Q)\theta$ for all θ in a neighborhood of the origin. Assuming $\operatorname{Im} Q \subseteq \operatorname{Im} R$, we can choose $H_c = -R^\dagger Q$, where $R^\dagger$ is the pseudo-inverse of R. Note that, by substituting into (9), we have

$$y(t) = RH_c\theta(t) + Q\theta(t) = 0, \quad \forall t \in \mathbb{R}_{\geq 0}.$$

Namely, the gradient is identically zero at all times. We also note that this is a particular feature of the parameter feedback problem, which can be achieved because $\theta(t)$ is assumed to be known when implementing (8). As we will see shortly below (cf. Section IV), $y(t) = 0$ can be achieved only as $t \to \infty$ for the gradient-feedback problem.

We conclude by noting that $y(t) = 0$ for all $t \in \mathbb{R}_{\geq 0}$ holds for all initial conditions $\theta(0)$ because the mapping $H_c(\theta)$ derived here zeros the gradient globally (and not only in some neighborhood of the origin). See also Theorem 3. □

IV. The dynamic gradient-feedback problem

In this section, we will build on the framework derived for the parameter-feedback problem to address the main objectives of this work: Problem 1.

A. Fundamental results

We begin with the following instrumental characterization.

Theorem 4 (Gradient-feedback algorithm characterization)**.** Suppose Assumptions 1, 2, and 3 hold, and assume that $F_c(z,y)$ and $G_c(z)$ are such that the equilibrium $z = z_\circ^\star$ of

$$\dot z(t) = F_c(z(t), \nabla_x f(G_c(z(t)), \theta(t))),$$

is exponentially stable. Then, the gradient-feedback optimization algorithm (5) asymptotically tracks a critical trajectory of (1) if and only if there exists a C^2 mapping $z = \sigma(\theta)$ with $\sigma(0) = z_\circ^\star$, defined in a neighborhood $\Theta_\circ \subset \Theta$ of the origin, that satisfies the following for all $\theta \in \Theta_\circ$:

$$\frac{\partial \sigma(\theta)}{\partial \theta} s(\theta) = F_c(\sigma(\theta), 0), \qquad (13a)$$

$$0 = \nabla_x f(G_c(\sigma(\theta)), \theta). \qquad (13b)$$

Proof. (Only if). We first prove that $\lim_{t\to\infty} y(t) = 0$ implies (13). The coupled dynamics (5) have the form:

$$\begin{aligned} \dot z &= (A_c + B_c R M) z + B_c Q \theta + \chi(x,\theta), \\ \dot \theta &= S\theta + \psi(\theta). \end{aligned} \qquad (14)$$

for some mappings $\chi(x,\theta)$ and $\psi(\theta)$ that vanish at the origin along with their first-order derivatives, and where Q and S are defined in (7) and

$$A_c = \left[\frac{\partial F_c}{\partial z}\right]_{(z,y)=(z_\circ^\star,0)}, \qquad B_c = \left[\frac{\partial F_c}{\partial y}\right]_{(z,y)=(z_\circ^\star,0)},$$
$$R = \left[\frac{\partial \nabla_x f}{\partial x}\right]_{(x,\theta)=(x_\circ^\star,0)}, \qquad M = \left[\frac{\partial G_c}{\partial z}\right]_{z=z_\circ^\star}.$$

By assumption, the eigenvalues of the matrix $A_c + B_c RM$ are in $\mathbb{C}^-$ and those of S on the imaginary axis. By Theorem 8, the system (14) has a center manifold at (0,0): the graph of a mapping $z=\sigma(\theta)$, with $\sigma(\theta)$ satisfying (see (31))

$$\frac{\partial\sigma(\theta)}{\partial\theta}s(\theta) = F_c(\sigma(\theta), \nabla_x f(G_c(\sigma(\theta)),\theta)).$$

By Assumption 3, no trajectory on this manifold converges to zero, so $\lim_{t\to\infty} y(t) = 0$ can hold only if (13b) holds, in which case the above equation reduces to (13a).

(If). We now prove that (13) implies $\lim_{t\to\infty} y(t)=0$. By Theorem 9, the center manifold $z=\sigma(\theta)$ is locally attractive; namely, $z(t)\to\sigma(\theta(t))$ as $t\to\infty$. Then, the fulfillment of (13b) guarantees that $y(t)\to 0$. □

The two conditions in (13) fully characterize the class of optimization algorithms that achieve asymptotic tracking of a critical trajectory. In words, (4) tracks a critical trajectory if and only if, for some mapping σ, the composite function $G_c\circ\sigma$ zeros the gradient locally (see (13b)), and the controller $F_c(z,y)$ is algebraically related to the exosystem $s(\theta)$ as given by (13a). Notice that, by Theorem 1, the former condition implies that

$$x(t) = G_c(\sigma(\theta)), \tag{15}$$

is a parameter-feedback optimization algorithm for (1).

Remark 3 (The internal model principle)**.** We interpret condition (13a) as the *internal model principle of time-varying optimization,* as it expresses the requirement that any optimization algorithm that achieves asymptotic tracking must include an internal model of the exosystem. Note that the use of a copy of the temporal variability of the optimization problem is explicit in the prediction-correction algorithm [10] (precisely, through the term $\nabla_{xt}f_0(x(t),t)$ – see Remark 1 for notation). □

It is important to note that, by Theorem 4, the exosystem state θ and that of the optimization z must be related, everywhere in $\Theta_\circ$, by the relationship:

$$z(t) = \sigma(\theta(t)). \tag{16}$$

Intuitively, (16) is interpreted as the existence of a change of coordinates between the state of the exosystem and that of the optimization (we discuss the invertibility properties of $\sigma(\theta)$ shortly below; see Section IV-B and Remark 4).

Remark 4. An important special case is obtained when σ is the identity operator on Θ; in this case, (13) simplifies to:

$$s(\theta) = F_c(\theta, 0),$$

which states that the controller vector field $F_c(z,y)$ must coincide with that of the exosystem $s(\theta)$ in $\Theta_\circ$. In this case, (16) gives $z(t)=\theta(t)$; namely, the controller state $z(t)$ and that of the exosystem $\theta(t)$ coincide in $\Theta_\circ$. □

While Theorem 4 provides a full characterization of all gradient-feedback algorithms that achieve tracking, it remains to address under what conditions on the loss $f(x,\theta)$ such an algorithm is guaranteed to exist. This question is addressed by the following result.

Theorem 5 (Existence of gradient-feedback algorithms)**.** Suppose Assumptions 1, 2, 3, and 4 hold. There exists a gradient-feedback optimization algorithm that solves Problem 1 if and only if there exists a mapping $H_c : \Theta \to \mathbb{R}^n$ that zeros the gradient locally. □

Proof. (Only if) By Theorem 4, there exists a mapping $z=\sigma(\theta)$ such that (13b) holds. Then, (10) holds immediately by letting $H_c(\theta) = G_c(\sigma(\theta))$.

(If) We will prove this claim by constructing a gradient-feedback algorithm that achieves $y(t)\to 0$ as $t\to\infty$.

First, notice that by Assumption 4, there exists a matrix L such that $S-LQ$ has eigenvalues in $\mathbb{C}^-$. Consider the algorithm (4) with $n_c = p$ and[3]

$$F_c(z,y) = s(z) + L(y - \nabla_x f(H_c(z),z)),$$
$$G_c(z) = H_c(z),$$

where $H_c(z)$ is as in (10). The Jacobian of $\dot z(t) = F_c(z(t),y(t))$ with respect to z is given by $S-LQ$ (notice that two terms of the form LRM with opposite sign cancel out in forming the Jacobian: one from the gradient term $\nabla_x f(H_c(z),z)$ and the other one from y). The claim thus follows by application of Theorem 4 with σ the identity operator on Θ. □

Interestingly, the conditions for existence of a gradient-feedback algorithm and those for existence of a parameter-feedback algorithm are identical. This is not surprising, since a parameter-feedback algorithm is assumed to have access to $\theta(t)$ while a gradient-feedback algorithm needs to measure $\theta(t)$ indirectly through $y(t)$. More precisely, as stated by (16), the dynamic state of the controller $z(t)$ acts as alternative representation (i.e., in different coordinates) of the exosystem state $\theta(t)$, while the exploration signal $x(t)$ acts as a parameter-feedback algorithm (see (15)).

The proof of Theorem 5 is constructive, as it provides a design procedure to construct $F_c(z,y)$ and $G_c(z)$ that constitute a

[3]Note that the second argument of the gradient is evaluated at the algorithm state z instead of the parameter vector θ as the latter is unknown.

gradient-feedback optimization algorithm. Such a procedure is presented in Algorithm 1, where a Luenberger observer is used to estimate the exosystem state $\theta(t)$ (see line 4), and a parameter feedback algorithm is then applied to the estimated exosystem state to regulate the gradient to zero (precisely, $G_c(z)$ is designed following the approach of Theorem 1 – see line 3).

Algorithm 1: Gradient-feedback algorithm design

Data: $s(\theta)$, $\nabla_x f(x,\theta)$, $H_c(\theta)$ satisfying (13), Jacobian matrices Q and S in (7)

1 $n_c \leftarrow n$;

2 $L \leftarrow$ any matrix such that $S - LQ$ is Hurwitz;

3 $G_c(z) \leftarrow H_c(z)$;

4 $F_c(z,y) \leftarrow s(z) + L(y - \nabla_x f(H_c(z), z))$;

Result: $F_c(z,y)$, $G_c(z)$, and n_c that solve Problem 1

Remark 5. Instead of a Luenberger observer, alternative dynamic observers could be considered in Line 4 of the algorithm to achieve different asymptotic or transient properties of the resulting gradient-feedback optimization algorithm; we leave the investigation of alternative state observer algorithms as the scope of future works. □

We illustrate the applicability of (13) on a quadratic problem in the following example.

Example 4. Consider the quadratic problem studied in Example 3, and assume that the exosystem follows the linear model $\dot{\theta} = S\theta$ for some matrix $S \in \mathbb{R}^{p\times p}$. According to Theorem 5, an optimization algorithm given by

$$\dot{z} = A_c z + B_c y, \qquad x = G_c z, \qquad y = Rx + Q\theta, \qquad (17)$$

where $A \in \mathbb{R}^{n_c\times n_c}$, $B_c \in \mathbb{R}^{n_c\times n}$, $G_c \in \mathbb{R}^{n\times n_c}$, achieves asymptotic tracking if and only if there exists a linear transformation $\Sigma \in \mathbb{R}^{n_c\times p}$ such that:

$$\Sigma S = A_c \Sigma, \qquad (18a)$$

$$0 = (RG_c\Sigma + Q), \qquad (18b)$$

for all θ in a neighborhood of the origin. When this condition holds, the application of Algorithm 1 gives

$$A_c = S, \qquad B_c = L, \qquad G_c = -R^\dagger Q,$$

where L is any matrix such that $S - LQ$ is Hurwitz; notice that this choice satisfies (18) with $\Sigma = I$. □

B. The optimization algorithm as a copy of the exosystem

As discussed after (16), the state of any optimization algorithm that achieves asymptotic tracking and that of the exosystem must be related by a change of coordinates. In the following, we show that $\sigma(\theta)$ is indeed an injective map, and it can thus be interpreted as a change of coordinates. We begin with the following result.

Proposition 6. Let the assumptions of Theorem 5 hold and let $F_c(z,y)$ and $G_c(z)$ be obtained by Algorithm 1. Then, there exists an injective map $\sigma(\theta)$ such that (16) holds. □

Proof. The claim follows by noting that $F_c(z,y)$ and $G_c(z)$ obtained by Algorithm 1 satisfy (13) with σ the identity operator on Θ. □

The interpretation of Proposition 6 is that, for any solution returned by Algorithm 1, $F_c(z,y)$ incorporates a copy of $s(\theta)$; precisely,

$$F_c(\sigma(\theta), 0) = s(\theta). \qquad (19)$$

The above feature is not a result of using Algorithm 1. Indeed, every gradient feedback algorithm has this feature, as shown next.

Proposition 7. Let the assumptions of Theorem 5 hold and assume that $S(\theta)$ is C^∞. For any $F_c(z,y)$ and $G_c(z)$ that achieve asymptotic tracking, there exists a neighborhood $\Theta_\circ \subset \Theta$ of the origin such that $\sigma(\theta)$ is injective in $\Theta_\circ$. □

Proof. By contradiction, assume that $F_c(z,y)$ and $G_c(z)$ satisfy (13), but $\sigma(\theta)$ is not injective at the origin; namely, there exists nonzero $\theta' \in \Theta_\circ$ such that $\sigma(\theta') = z_0^\star$. From (13b), we have

$$0 = \nabla_x f(G_c(\sigma(\theta')), \theta').$$

Moreover, from Theorem 4,

$$\frac{\partial \sigma(\theta')}{\partial \theta} s(\theta') = F_c(\sigma(\theta'), 0) = 0. \qquad (20)$$

For C^∞ vector fields $h_1(x)$ and $h_2(x)$, we let

$$L_{h_1}(h_2)(x) = \frac{\partial h_2(x)}{\partial x} h_1(x).$$

By application of (20), we have

$$0 = L_s(\nabla_x f)(x,\theta') = L_s(L_s(\nabla_x f))(x,\theta') = \dots.$$

Hence, the matrix

$$\begin{bmatrix} L_s(\nabla_x f)(x,\theta) \\ L_s(L_s(\nabla_x f))(x,\theta) \\ \vdots \end{bmatrix},$$

is not invertible at $\theta = \theta'$. By [35, Thm. 3.13], the system (2) is not weakly observable, thus violating Assumption 4. □

V. Fidelity of the internal model and tracking accuracy

In this section, we illustrate through examples how an imprecise knowledge of the internal model impacts the tracking accuracy. For tractability, we will focus on the quadratic problem (12) (see Examples 3 and 4). Suppose that an imprecise knowledge of the internal model is available; namely, there exists $\Delta \in \mathbb{R}^{n_c\times p}$ such that (18) is modified to:

$$\Sigma S + \Delta = A_c \Sigma, \qquad (21a)$$

$$0 = (RG_c\Sigma + Q). \qquad (21b)$$

To analyze the asymptotic behavior of $y(t)$, it is useful to define the auxiliary variable $\tilde{z}(t) = z(t) - \Sigma\theta(t)$. Using (21) and (17), the dynamics of $\tilde{z}$ follow the model:

$$\dot{\tilde{z}} = (A_c + B_c R G_c)\tilde{z} + \Delta\theta. \tag{22a}$$

The corresponding gradient signal in terms of $\tilde{z}$ is:

$$y = R G_c \tilde{z}. \tag{22b}$$

Assuming that $(A_c + B_c R G_c)$ is Hurwitz stable (see Theorem 4), the Final Value Theorem gives:

$$\begin{aligned} y_\infty &:= \lim_{t\to\infty} y(t) \\ &= \lim_{s\to 0} sRG_c(sI - A_c - B_cRG_c)^{-1}\Delta(sI - S)^{-1}\theta(0). \end{aligned} \tag{23}$$

As a first illustrative scenario, suppose $S = 0$; namely, that the exosystem states are constants at all times. Notice that, since $S = 0$, from (18a), A_c incorporates an internal model of S if and only if $A_c\Sigma = 0$; precisely, if and only if

$$\dim(\ker(A_c)) \geq p.$$

In this case, assuming that the perturbed A_c is such that $(A_c + B_cRG_c)$ remains Hurwitz stable, from (23), we have:

$$y_\infty = -RG_c(A_c + B_cRG_c)^{-1}\Delta\theta(0).$$

As expected, when $\Delta = 0$, A_c incorporates an exact internal model, and $y_\infty = 0$ for any $\theta(0)$. On the other hand, when $\Delta \neq 0$,

$$\|y_\infty\| \leq \|RG_c(A_c + B_cRG_c)^{-1}\| \|\theta(0)\| \|\Delta\|,$$

namely, bounded errors in Δ result in a bounded y_∞.

Interestingly, analogous continuity properties may not hold when the Poisson stability assumption is dropped. Consider an instance of (12) with $n = p = 2$, $R = Q = I$, and

$$S = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}.$$

Suppose the algorithm (17) is used, with $G_c = I, B_c = I$,

$$A_c = \begin{bmatrix} -\epsilon_1 & 1 \\ 0 & -\epsilon_2 \end{bmatrix},$$

where $0 \leq \epsilon_1 < 1$, $0 \leq \epsilon_2 < 1$. To avoid trivial cases, we will assume $\theta_1(0)$ and $\theta_2(0)$ are nonzero. When $\epsilon_1 = \epsilon_2 = 0$, (21) holds with $\Sigma = -I$ and $\Delta = 0$, in which case $y_\infty = 0$. On the other hand, when $\epsilon_1, \epsilon_2 \neq 0$, from (23):

$$y_\infty = \lim_{s\to 0} \begin{bmatrix} \frac{\epsilon_1}{s-\alpha_1} & \frac{\epsilon_1}{s(s-\alpha_1)} + \frac{\epsilon_2}{(s-\alpha_1)(s-\alpha_2)} \\ 0 & \frac{\epsilon_2}{s-\alpha_2}, \end{bmatrix} \begin{bmatrix} \theta_1(0) \\ \theta_2(0) \end{bmatrix}$$

where $\alpha_i = 1 - \epsilon_i$ for $i \in \{1, 2\}$. Since $\epsilon_1 \neq 0$ and $\theta_2(0) \neq 0$, we conclude that $|y(t)| \to \infty$ as $t \to \infty$, i.e., asymptotic tracking is not attained.

In conclusion, the tracking accuracy depends on the fidelity of the internal model as well as the asymptotic behavior of the exosystem.

VI. Extensions to constrained problems

We can extend our results to study constrained time-varying optimization problems by searching for a stationary point of the Lagrangian.

Consider the equality-constrained problem

$$\begin{aligned} &\text{minimize} \quad f(x, \theta(t)) \\ &\text{subject to} \quad h_i(x, \theta(t)) = 0, \quad i = 1, \dots, m \end{aligned}$$

where the constraint functions $h_i(x, \theta(t))$ depend on the parameter vector. The associated Lagrangian function is

$$L(x, \lambda, \theta(t)) = f(x, \theta(t)) + \sum_{i=1}^{m} \lambda_i h_i(x, \theta(t))$$

where λ_i is the Lagrange multiplier associated with the i^{th} equality constraint. A pair $(x, \lambda) \in \mathbb{R}^n \times \mathbb{R}^m$ is said to be a saddle-point of the Lagrangian if

$$L(x, \bar{\lambda}, \theta(t)) \leq L(x, \lambda, \theta(t)) \leq L(\bar{x}, \lambda, \theta(t))$$

for all pairs $(\bar{x}, \bar{\lambda}) \in \mathbb{R}^n \times \mathbb{R}^m$. For any such saddle-point, x is primal optimal, λ is dual optimal, and the optimal duality gap is zero. Moreover, the gradient of the Lagrangian (assuming it exists) is zero at any saddle-point. It follows from the derivations in the previous sections that the gradient-feedback and parameter-feedback algorithms can be directly applied to seek a stationary point of the Lagrangian function by replacing (in (1)) the variable x with the extended decision variable $\tilde{x} = (x, \lambda)$ and by letting $f(\tilde{x}, \theta) = L(x, \lambda, \theta)$. Notice that, if the critical point computed by (4a) is also a saddle-point, then it is also a solution to the equality-constrained problem; see [36, Ch. 5].

VII. Simulation results

In this section, we illustrate our results and optimization design method through a set of numerical simulations.

A. Optimization design for quadratic costs

We begin by numerically investigating the quadratic instance of (1) with linear temporal variability, discussed previously in Example 3. With dimensions $n = p = 4$, we chose the matrix $R \in \mathbb{S}^4$ with random entries such that its eigenvalues are uniformly distributed in the open real interval $(0, 1)$, and we set $Q = I \in \mathbb{R}^{4\times 4}$. We let the exosystem be $\dot{\theta}(t) = S\theta(t)$, where $S \in \mathbb{R}^{4\times 4}$ is given by $S = \tilde{S} - \tilde{S}^\mathsf{T}$, and $\tilde{S}$ is a matrix with random entries uniformly distributed in the open real interval $(0, 1)$. Notice that this choice ensures that the eigenvalues of S are on the imaginary axis. We chose $H_c(\theta) = H_c\theta$ with $H_c = -R^{-1}Q$, which ensures (10) holds (see Example 3). We applied Algorithm 1, choosing L such the eigenvalues of $S - LQ$ are uniformly distributed in the real interval $(-2, -1)$, letting $G_c = -R^{-1}Q$, and

$$\begin{aligned} F_c(z, y) &= s(z) + L(y - \nabla_x f(H_c(z), z)), \\ &= Sz + L(y - (RG_cz + Qz)), \\ &= Sz + Ly. \end{aligned}$$

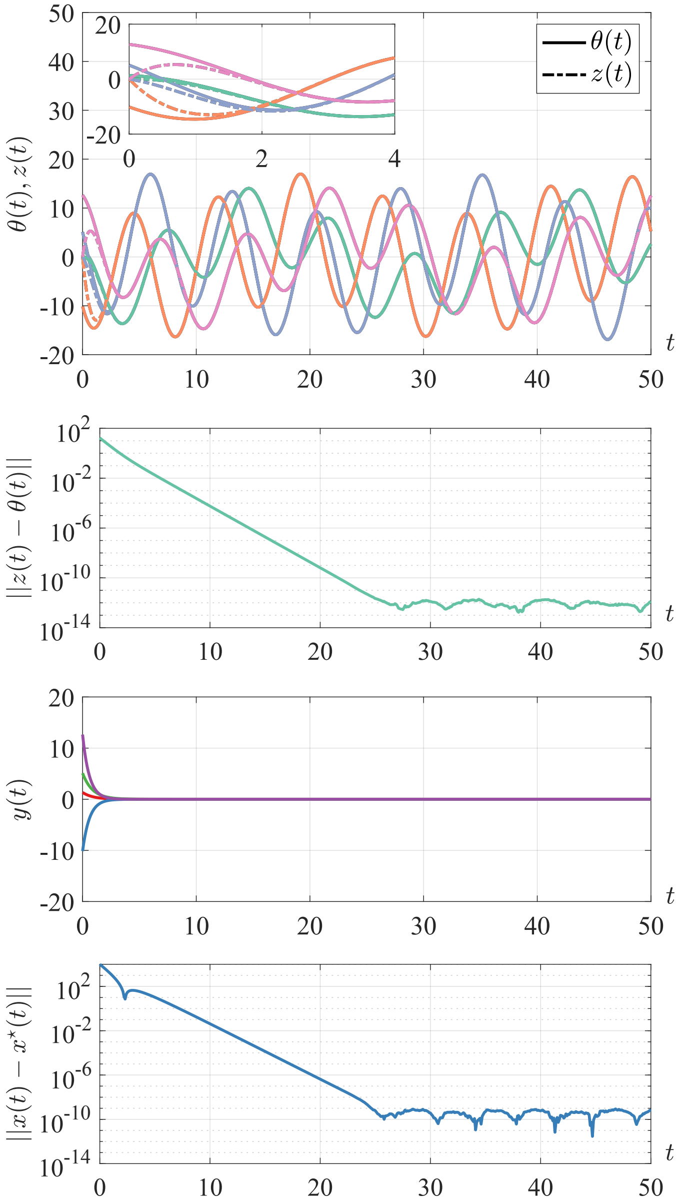

Fig. 3: Simulation results illustrating the performance of an optimization algorithm synthesized using Algorithm 1 for the quadratic instance (12) of (1). See Example 3 and Section VII-A for a discussion. (Top) Illustration of the temporal variability of the parameter $\theta(t)$ and of $z(t)$. (Second from top) $z(t)$ is an estimator for $\theta(t)$, and thus $z(t) \to \theta(t)$ as $t \to \infty$. (Third from top) The proposed control algorithm is successful in regulating the gradient feedback signal $y(t)$ to zero asymptotically. (Bottom) Illustration that $x(t) \to x^\star(t)$ as $t \to \infty$.

The simulation results are presented in Figure 3. It follows from our choice of $F_c(z, y)$, that the optimization state $z(t)$ converges to the parameter vector $\theta(t)$ (see the top two plots in Fig. 3). Moreover, it follows from our choice of $G_c(z)$ that $y(t) \to 0$ and thus $x(t) \to x^\star(t)$ as $t \to \infty$ (see the bottom two plots in Fig. 3).

B. Application to solve the dynamic traffic assignment problem in transportation

We next illustrate the applicability of the framework in solving the dynamic traffic assignment problem in roadway transportation [37]; intuitively, the objective is to decide how traffic flows split among the available paths of a network to minimize the drivers' travel time to destination. We model a roadway transportation network using a static flow model [37], described by a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with edges $i \in \mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ (modeling traffic roads) and nodes $\mathcal{V}$ (modeling traffic junctions). For $i \in \mathcal{E}$, we denote by $i^+ \in \mathcal{V}$ and $i^- \in \mathcal{V}$ its origin and destination nodes, respectively. We assume that an exogenous, time-varying, inflow of traffic $\theta(t)$ enters the network at a certain origin node, denoted by $o \in \mathcal{V}$, and exits at a certain destination node, denoted by $d \in \mathcal{V}$; for simplicity, we assume that there is only one origin-destination pair, but this is without loss of generality [37]. We describe the network state using a vector $x \in \mathbb{R}^{|\mathcal{E}|}_{\geq 0}$ (where $|\mathcal{E}|$ denotes the number of edges) whose entries x_i describe the amount of inflow θ routed through road i. To each link i, we associate a function $\ell_i(x_i)$ describing the latency (or travel time) of road i. For our simulations, we consider the network topology in Fig. 4, and choose the latency functions as follows:

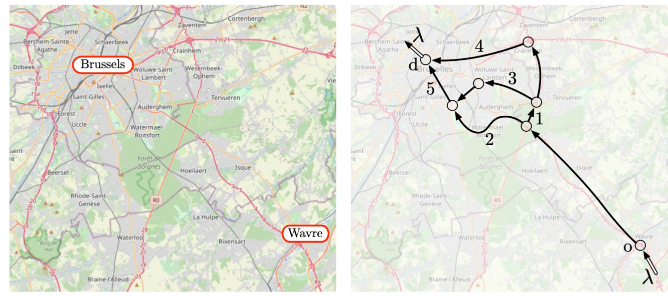

(a) (b)

Fig. 4: (Left) Areal view of the highway system between the cities of Wavre and Brussels, Belgium. (Right) Graph utilized to model the portion of traffic network of interest.

$$\begin{aligned} &\ell_1(x_1) = x_1, &&\ell_2(x_2) = 10x_2, &&\ell_3(x_3) = x_3, \\ &\ell_4(x_4) = 5x_4, &&\ell_5(x_5) = x_5. \end{aligned}$$

According to Wardrop's first principle [37, pp. 31], transportation networks operate at a condition where travelers select their path to minimize their travel time to destination. Mathematically, a Wardrop's equilibrium is the optimizer of the following optimization problem:

$$\begin{aligned} \min_{x \in \mathbb{R}^{|\mathcal{E}|}} \quad & \sum_{i \in \mathcal{E}} \int_0^{x_i} \ell_i(s) ds \\ \text{subject to:} \quad & \sum_{j \in \mathcal{E}: j^- = v} x_j - \sum_{j \in \mathcal{E}: j^+ = v} x_j = \delta_v(\theta(t)), \ \forall v \in \mathcal{V}, \\ & x_i \geq 0, \ \forall i \in \mathcal{E}, \end{aligned} \tag{24}$$

where $\delta_v(\theta)$, $v \in \mathcal{V}$, is defined as:

$$\delta_v(\theta) = \begin{cases} \theta, & \text{if } v = o, \\ -\theta, & \text{if } v = d, \\ 0, & \text{otherwise.} \end{cases}$$

The loss function in (24) is used to model travelers who will switch to a different path if it has shorter travel time to destination, while the first constraint in (24) describes the network topology, namely, that traffic flows are conserved at each node. Notice that (24) is a time-varying optimization problem, where the temporal variability originates from the dependence of the constraint on $\theta(t)$, which describes the inflow of vehicles at the origin and outflow at the destination, measured in vehicles per hour. For our simulations, we assume that the network inflow is sinusoidal:

$$\theta(t) = \theta_0 - \theta_1 \cos(\omega_1 t + \phi_1) - \theta_2 \cos(\omega_2 t + \phi_2), \qquad (25)$$

where $\theta_0, \theta_1, \theta_2, \omega_1, \omega_2, \phi_1, \phi_2 \in \mathbb{R}_{>0}$, satisfy $\theta_0 > \theta_1$, $\theta_0 > \theta_2$, and $\omega_2 > \omega_1$. The model (25) states that the network inflow is the sum of a constant term, θ_0, a slowly-varying sinusoid with angular frequency ω_1 and a quickly-varying sinusoid with angular frequency ω_2. The low-frequency sinusoid is used here to describe slowly-varying (e.g., hourly) traffic demands, while the high-frequency sinusoid is used to model sudden (e.g., at the minute-level) variations in traffic demand. For our simulations, we let $\theta_0 = 3$ veh/h, $\theta_1 = 1$ veh/h, $\theta_2 = 0.1$ veh/h, $\omega_1 = 0.1$ rad/hour, $\omega_2 = \sqrt{50}$ rad/hour, and $\phi_1 = \phi_2 = 0$, see Fig. 5 (top).

We applied Algorithm 1 to derive an optimization algorithm to solve the traffic assignment problem (24). For the synthesis, we utilized the internal model $s(z) = Sz$, where

$$S = \begin{bmatrix} 0 & 1 & 0 & 0 & 0 \\ -\omega_1^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & -\omega_2^2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

Notice that knowledge of $\theta_0, \theta_1, \theta_2, \phi_1, \phi_2$ is not required to synthesize the optimization algorithm – only the frequencies ω_1, ω_2 are required to be known. To seek a solution to the constrained problem (24), consider the Lagrangian function:

$$L(x, \lambda, \theta(t)) := \sum_{i \in \mathcal{E}} \int_0^{x_i} \ell_i(s) ds + \sum_{v \in \mathcal{V}} \lambda_v \left(\delta_v(\theta(t)) - \sum_{\substack{j \in \mathcal{E}: \\ j^- = v}} x_j + \sum_{\substack{j \in \mathcal{E}: \\ j^+ = v}} x_j \right),$$

where $\lambda := (\lambda_1, \dots, \lambda_{|\mathcal{V}|})$ is the vector of Lagrange multipliers. We applied Algorithm 1 to the optimization problem (1) with $f(\tilde{x}, \theta) = L(x, \lambda, \theta)$ (see Section VI); the inequality constraints in (24) have been accounted for by projecting $x(t)$ onto the feasible set. Here, matrix L has been chosen so that the eigenvalues of $S - LQ$ are uniformly distributed in the open real interval $(-1, -2)$. Notice that, since the latencies are strictly increasing, the Lagrangian is strongly convex-strongly concave [37], and thus the problem admits a unique critical point that is a saddle point. It follows that our algorithm is guaranteed to converge to a minimizer of (24). Simulation for this problem are presented in Fig. 5. The simulation shows that $z(t) \to \theta(t)$ as $t \to \infty$ (see Fig. 5-second figure) and $y(t) \to 0$, which implies that the algorithm successfully

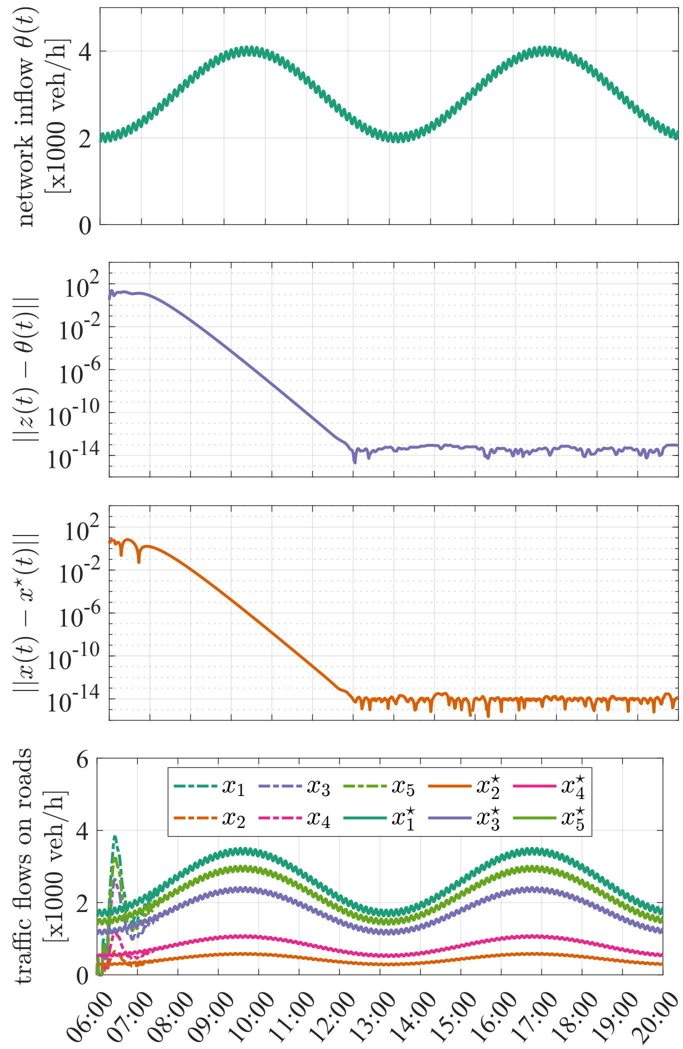

Fig. 5: Simulation results illustrating the performance of the optimization algorithm synthesized using Algorithm 1 to solve the dynamic traffic assignment problem (24). See Section VII-B for a discussion. The bottom figure show that, from a zero initial condition, the algorithm is capable of computing a Wardrop equilibrium in about 1.5 hours, and then of tracking this equilibrium. Notice that, a faster rate of convergence could be obtained by shifting the eigenvalues of $S - LQ$.

computes the critical point (see Fig. 5-third figure). Finally, the bottom figure of Fig. 5 illustrates the rate of convergence of the algorithm, which, as expected, is governed by the placement of the eigenvalues of the observer.

VIII. Conclusions

We showed that the problem of designing optimization algorithms for time-varying optimization problems can be reconducted to the output regulation problem of nonlinear multivariable systems. This connection allowed us to prove the internal model principle of time-varying optimization, which states that asymptotic tracking can be achieved only if the optimization algorithm incorporates a reduplicated model of the temporal variability of the problem. Further, we showed that asymptotic tracking can be achieved under more relaxed assumptions (namely, assumption (A2)) than what is normally imposed in the literature (see assumption (A3)). Moreover, our algorithm structure is novel in the literature, and it relies on the use of an observer for the temporal variability of the problem. This work opens the opportunity for several directions of future work, including an investigation at discrete-time and possibly the application in feedback optimization.

Appendix

We now summarize relevant facts in center manifold theory from [18]; see also [38]. Consider the nonlinear system:

$$\dot{x} = f(x) \tag{26}$$

where f is a C^k vector field defined on an open subset U of $\mathbb{R}^n$, and let $x_\circ \in U$ be an equilibrium point for f, i.e., $f(x_\circ) = 0$. Without loss of generality, suppose $x^\circ = 0$. Let $F = \left[\frac{\partial f}{\partial x}\right]_{x=0}$, denote the Jacobian matrix of f at $x = 0$. Suppose the matrix F has n° eigenvalues with zero real part, n^- eigenvalues with negative real part, and n^+ eigenvalues with positive real part. Let E^-, E°, and E^+ be the (generalized) real eigenspaces of F associated with eigenvalues of F lying on the open left half plane, the imaginary axis, and the open right half plane, respectively. Note that E°, E^-, E^+ have dimension n°, n^-, n^+, respectively and that each of these spaces is invariant under the flow of $\dot{x} = Fx$. If the linear mapping F is viewed as a representation of the differential (at $x = 0$) of the nonlinear mapping f, its domain is the tangent space T_0U to U at $x = 0$, and the three subspaces in question can be viewed as subspaces of T_0U satisfying $T_0U = E^\circ \oplus E^- \oplus E^+$. We refer to [39, Sec. A.II] for a precise definition of C^k manifolds; loosely speaking, a set $S \subset U$ is a C^k manifold it can be locally represented as the graph of a C^k function.

Definition 3 (Locally invariant manifold)**.** A C^k manifold S of U is said to be locally invariant for (26) if, for each $x_\circ \in S$, there exists $t_1 < 0 < t_2$ such that the integral curve $x(t)$ of (26) satisfying $x(0) = x_\circ$ is such that $x(t) \in S$ for all $t \in (t_1, t_2)$. □

Intuitively, by letting $x = (y, \theta)$ and expressing (26) as:

$$\dot{y} = f_y(\theta, y), \qquad \dot{\theta} = f_\theta(\theta, y), \tag{27}$$

a curve $y = \pi(\theta)$ is an invariant manifold for (27) if the solution of (27) with $\theta(0) = \theta_\circ$ and $y(0) = \pi(\theta_\circ)$ lies on the curve $y = \pi(\theta)$ for t in a neighborhood of 0. The notion of invariant manifold is useful as, under certain assumptions, it allow us to reduce the analysis of (26) to the study of a reduced system in the variable θ only. The remainder of this section is devoted to formalizing this fact.

Definition 4 (Center manifold)**.** Let $x = 0$ be an equilibrium of (26). A manifold S, passing through $x = 0$, is said to be a center manifold for (26) at $x = 0$ if it is locally invariant and the tangent space to S at 0 is exactly E°. □

Returning to the decomposition (27), intuitively, the invariant manifold $y = \pi(\theta)$ is said to be a center manifold when all orbits of y decay to zero and those of θ neither decay nor grow exponentially.

In what follows, we will assume that all eigenvalues of F have nonpositive real part, i.e.,

$$n^+ = 0. \tag{28}$$

When (28) holds, it is always possible choose coordinates in U such that (26) reads as:

$$\dot{y} = Ay + g(y, \theta) \tag{29a}$$

$$\dot{\theta} = B\theta + h(y, \theta) \tag{29b}$$

where A is an $n^- \times n^-$ matrix having all eigenvalues with negative real part, B is an $n^\circ \times n^\circ$ matrix having all eigenvalues with zero real part, and the functions g and h are C^k functions vanishing at $(y, \theta) = (0, 0)$, together with all their first-order derivatives. Because of their equivalence, any conclusion drawn for (29) will apply also to (26). The following result ensures the existence of a center manifold.

Theorem 8 (Center manifold existence theorem)**.** Assume that (28) holds. There exist a neighborhood $V \subset \mathbb{R}^{n^0}$ of 0 and a class C^{k-1} mapping $\pi : V \to \mathbb{R}^{n^-}$ such that the set

$$S = \{(y, \theta) \in \mathbb{R}^{n^-} \times V : y = \pi(\theta)\},$$

is a center manifold for (29). □

Some important observations are in order. By definition, a center manifold for (29) passes through $(0, 0)$ and is tangent to the subset of points whose y coordinate is zero. Namely,

$$\pi(0) = 0 \qquad \text{and} \qquad \frac{\partial \pi}{\partial \theta}(0) = 0. \tag{30}$$

Moreover, this manifold is locally invariant for (29): this imposes on the mapping π the constraint:

$$\frac{\partial \pi}{\partial \theta}(B\theta + h(\pi(\theta), \theta)) = A\pi(\theta) + g(\pi(\theta), \theta), \tag{31}$$

as deduced by differentiating with respect to time any solution $(y(t), \theta(t))$ of (29) on the manifold $y(t) = \pi(\theta(t))$. In other words, any center manifold for (29) can equivalently be described as the graph of a mapping $y = \pi(\theta)$ satisfying the partial differential equation (31), with the constraints specified by (30).

Remark 6. Theorem 8 shows existence but not uniqueness of a center manifold. Moreover, (i) if g and h are $C^k, k \in \mathbb{N}_{>0}$, then (29) admits a C^{k-1} center manifold; (ii) if g and h are C^∞ functions, then (29) has a C^k center manifold for any finite k, but not necessarily a C^∞ center manifold. □

The next result shows that any y-trajectory of (29), starting sufficiently close to the origin, converges, as time tends to infinity, to a trajectory that belongs to the center manifold.

Theorem 9. Assume that (28) holds and suppose $y = \pi(\theta)$ is a center manifold for (29) at $(0, 0)$. Let $(y(t), \theta(t))$ be a solution of (29). There exists a neighborhood U° of $(0, 0)$ and real numbers $M > 0$ and $K > 0$ such that, if $(y(0), \theta(0)) \in U^\circ$, then for all $t \geq 0$,

$$\|y(t) - \pi(\theta(t))\| \leq Me^{-Kt}\|y(0) - \pi(\theta(0))\|. \qquad \square$$

From the above discussion, any trajectory of (29) starting at a point $y^\circ = \pi(\theta^\circ)$ of a center manifold satisfies:

$$y(t) = \pi(\zeta(t)), \qquad \theta(t) = \zeta(t),$$

where $\zeta(t)$ is any solution of

$$\dot{\zeta} = B\zeta + h(\pi(\zeta), \zeta), \tag{32}$$

satisfying the initial condition $\zeta(0) = \theta_\circ$. This decomposition allows us to predict the asymptotic behavior of (29) by studying the asymptotic behavior of a reduced-order system, namely, (32). This is formalized in the following result.

Theorem 10. Suppose $\zeta = 0$ is a stable (respectively, asymptotically stable, unstable) equilibrium of (32). Then $(y, \theta) = (0, 0)$ is a stable (respectively, asymptotically stable, unstable) equilibrium of (29). $\square$

References

[1] Y. Chen and Y. Zhou, "Machine learning based decision making for time varying systems: Parameter estimation and performance optimization," *Knowledge-Based Systems*, vol. 190, p. 105479, 2020.

[2] A. Rakhlin and K. Sridharan, "Online learning with predictable sequences," in *Conference on Learning Theory*, 2013, pp. 993–1019.

[3] F. Y. Jakubiec and A. Ribeiro, "D-map: Distributed maximum a posteriori probability estimation of dynamic systems," *IEEE Transactions on Signal Processing*, vol. 61, no. 2, pp. 450–466, 2012.

[4] A. Davydov, V. Centorrino, A. Gokhale, G. Russo, and F. Bullo, "Contracting dynamics for time-varying convex optimization," *arXiv preprint*, 2023, arXiv:2305.15595.

[5] B. Chachuat, B. Srinivasan, and D. Bonvin, "Adaptation strategies for real-time optimization," *Computers & Chemical Engineering*, vol. 33, no. 10, pp. 1557–1567, 2009.

[6] V. M. Zavala and M. Anitescu, "Real-time nonlinear optimization as a generalized equation," *SIAM Journal on Control and Optimization*, vol. 48, no. 8, pp. 5444–5467, 2010.

[7] U. Helmke and J. B. Moore, *Optimization and dynamical systems*. Springer Science & Business Media, 2012.

[8] A. Cherukuri, E. Mallada, and J. Cortés, "Asymptotic convergence of constrained primal–dual dynamics," *Systems & Control Letters*, vol. 87, pp. 10–15, 2016.

[9] G. Bianchin, J. Cortés, J. I. Poveda, and E. Dall'Anese, "Time-varying optimization of LTI systems via projected primal-dual gradient flows," *IEEE Transactions on Control of Network Systems*, vol. 9, no. 1, pp. 474–486, Mar. 2022. doi: 10.1109/TCNS.2021.3112762.

[10] Y. Zhao and M. N. S. Swamy, "A novel technique for tracking time-varying minimum and its applications," in *IEEE Canadian Conference on Electrical and Computer Engineering*, vol. 2, 1998, pp. 910–913.

[11] M. Fazlyab, S. Paternain, V. M. Preciado, and A. Ribeiro, "Prediction-correction interior-point method for time-varying convex optimization," *IEEE Transactions on Automatic Control*, vol. 63, no. 7, pp. 1973–1986, 2017.

[12] N. Bastianello, R. Carli, and S. Zampieri, "Internal model-based online optimization," *IEEE Transactions on Automatic Control*, vol. 69, no. 1, pp. 689–696, 2024. doi: 10.1109/TAC.2023.3297504.

[13] A. Simonetto, E. Dall'Anese, S. Paternain, G. Leus, and G. B. Giannakis, "Time-varying convex optimization: Time-structured algorithms and applications," *Proceedings of the IEEE*, vol. 108, no. 11, pp. 2032–2048, 2020.

[14] A. Simonetto, E. Dall'Anese, J. Monteil, and A. Bernstein, "Personalized optimization with user's feedback," *Automatica*, vol. 131, p. 109767, 2021.

[15] B. A. Francis and W. M. Wonham, "The internal model principle of control theory," *Automatica*, vol. 12, no. 5, pp. 457–465, 1976.

[16] J. Hepburn and W. Wonham, "Error feedback and internal models on differentiable manifolds," *IEEE Transactions on Automatic Control*, vol. 29, no. 5, pp. 397–403, 1984.

[17] B. A. Francis, "The linear multivariable regulator problem," *SIAM Journal on Control and Optimization*, vol. 15, no. 3, pp. 486–505, 1977.

[18] J. Carr, *Applications of centre manifold theory*. Springer-Verlag, 1981, vol. 35. doi: 10.1007/978-1-4612-5929-9.

[19] E. C. Hall and R. M. Willett, "Online convex optimization in dynamic environments," *IEEE Journal of Selected Topics in Signal Processing*, vol. 9, no. 4, pp. 647–662, 2015.

[20] S. M. Fosson, "Centralized and distributed online learning for sparse time-varying optimization," *IEEE Transactions on Automatic Control*, vol. 66, no. 6, pp. 2542–2557, 2020.

[21] S. Shalev-Shwartz, "Online learning and online convex optimization," *Foundations and Trends® in Machine Learning*, vol. 4, no. 2, pp. 107–194, 2012.

[22] R. Dixit, A. S. Bedi, R. Tripathi, and K. Rajawat, "Online learning with inexact proximal online gradient descent algorithms," *IEEE Transactions on Signal Processing*, vol. 67, no. 5, pp. 1338–1352, 2019.

[23] M. Zinkevich, "Online convex programming and generalized infinitesimal gradient ascent," in *International conference on machine learning*, 2003, pp. 928–936.

[24] E. Hazan, A. Agarwal, and S. Kale, "Logarithmic regret algorithms for online convex optimization," *Machine Learning*, vol. 69, no. 2, pp. 169–192, 2007.

[25] E. Hazan, "Introduction to online convex optimization," *Foundations and Trends in Optimization*, vol. 2, no. 3-4, pp. 157–325, 2016.

[26] Y. Li, G. Qu, and N. Li, "Using predictions in online optimization with switching costs: A fast algorithm and a fundamental limit," in *American Control Conference*, 2018, pp. 3008–3013.

[27] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan, "Online optimization: Competing with dynamic comparators," in *Artificial Intelligence and Statistics*, 2015, pp. 398–406.

[28] O. Dekel, N. Haghtalab, and P. Jaillet, "Online learning with a hint," *Advances in Neural Information Processing Systems*, vol. 30, 2017.

[29] M. Marchi, J. Bunton, J. P. Silvestre, and P. Tabuada, "A framework for time-varying optimization via derivative estimation," *arXiv preprint*, 2024, arXiv:2403.19088.

[30] A. Hauswirth, Z. He, S. Bolognani, G. Hug, and F. Dörfler, "Optimization algorithms as robust feedback controllers," *Annual Reviews in Control*, vol. 57, p. 100941, 2024.

[31] A. Simonetto and E. Dall'Anese, "Prediction-correction algorithms for time-varying constrained optimization," *IEEE Transactions on Signal Processing*, vol. 65, no. 20, pp. 5481–5494, 2017.

[32] W. Rudin, *Principles of Mathematical Analysis*, 3rd ed., ser. International Series in Pure and Applied Mathematics. McGraw-Hill, 1976. ISBN 0-07-085613-3

[33] H. D. Nguyen and F. Forbes, "Global implicit function theorems and the online expectation–maximisation algorithm," *Australian & New Zealand Journal of Statistics*, vol. 64, no. 2, pp. 255–281, 2022.

[34] A. V. Arutyunov and S. E. Zhukovskiy, “Application of methods of ordinary differential equations to global inverse function theorems,” *Differential equations*, vol. 55, pp. 437–448, 2019.

[35] R. Hermann and A. Krener, “Nonlinear controllability and observability,” *IEEE Transactions on Automatic Control*, vol. 22, no. 5, pp. 728–740, 1977.

[36] S. Boyd and L. Vandenberghe, *Convex Optimization*. New York, NY, USA: Cambridge University Press, 2004. ISBN 0521833787

[37] M. Patriksson, *The traffic assignment problem: models and methods*. Mineola, NY: Dover Publications, 2015. ISBN 0486787907

[38] S. Wiggins, *Ordinary Differential Equations: A Dynamical Point of View*. World Scientific, 2023.

[39] M. C. Irwin, *Smooth dynamical systems*. World Scientific, 2001, vol. 17.