$\mathcal{L}_1$ Adaptive Control

Naira Hovakimyan
Professor, Department of Mechanical Science and Engineering
University of Illinois at Urbana-Champaign, Urbana, Illinois 61801, USA
e-mail: nhovakim@illinois.edu

The objective of this document is to provide a technical response to a widely circulated article [1], co-authored by P. Ioannou, K. Narendra, A. Annaswamy, S. Jafari, L. Rudd, R. Ortega, and J. Boskovic, in which the authors "have come together" to express "with one voice" their "growing concern" about the claims on performance and robustness guarantees of $\mathcal{L}_1$ adaptive control. The article has been submitted to IEEE Transactions on Automatic Control, and was also sent to funding agencies, government, industry, as well as academic communities. I was invited to provide a technical review for the article by IEEE Transactions on Automatic Control, which I formally declined because of two reasons. First, the document violates the objective standards for scholarly publications, and it seems to be written in an ill-spirited tone. Second, it is so wrong that it is clear that the authors did not put any effort to follow up on our work and rushed to create a document that shows complete lack of understanding of the main claims of the theory. Instead of providing a formal review, I composed this document to state –once more– the main properties of $\mathcal{L}_1$ adaptive control and present the technical mistakes of [1] for the community at large.

The first paper on $\mathcal{L}_1$ adaptive control was submitted to IEEE CDC in 2005, which was declined. It later appeared in ACC 2006 as a two-part paper, [2], [3]. It was later followed by a series of articles, culminating in the book, [4]. The development of $\mathcal{L}_1$ adaptive control was motivated by some of the apparent issues that model reference adaptive control (MRAC) has been facing for a number of years, [5]. Although every attempt has been made from our end to explain the novelty and the contribution of $\mathcal{L}_1$ adaptive control over the years in various papers and presentations, including [4], [6], [7], it seems that some of its benefits have ***not*** been well understood by several members of our community.

The key difference between MRAC and $\mathcal{L}_1$ adaptive control is in the **problem formulation**, which consequently leads to significant **architectural differences**. The wrong conception in [1] is that $\mathcal{L}_1$ adaptive control is perceived like an input filtered MRAC, which is ***not*** the case. While the $\mathcal{L}_1$ adaptive control architectures *do have* a filtering structure, there are **two key points** to be kept in mind: *i*) this filtering structure stands in a very particular point of the control architecture, ensuring that the estimation loop is decoupled from the control loop; and *ii*) the synthesis of this filtering structure depends on the class of systems considered and, in general, does not correspond to the mere insertion of a low-pass filter. These two points have to be carefully addressed with every particular system, dependent upon the presence of unmodeled dynamics, disturbances, noise, etc.

To make this exposition clear, this document includes three main parts: *i*) a description of the problem formulation, key ideas, and main properties of $\mathcal{L}_1$ adaptive control; *ii*) a brief overview of the current challenges of the theory; and *iii*) a discussion of the claims in [1], which –as it will become evident– are unsubstantiated and, for the most part, plainly wrong. In particular, we note that Section II-C analyzes in detail the simulation example from [1] with a **correct** implementation of an $\mathcal{L}_1$ adaptive controller.

I. Problem Formulation

The essential difference between MRAC and $\mathcal{L}_1$ adaptive control is the **problem formulation**. To articulate this difference, let's look into the following class of systems:

$$\begin{aligned} \dot{x}(t) &= A_m x(t) + b(u(t) - \theta^\top(t)x(t))\,, \qquad x(0) = x_0\,, \\ y(t) &= c^\top x(t)\,, \end{aligned} \tag{1}$$

where $\theta(t)$ is the vector of unknown parameters, confined to some *a priori* known compact set; A_m is a Hurwitz matrix specifying the desired closed-loop performance; b, c are known vectors; while x, u, and y are the state, the input, and the output of the system (of appropriate dimensions). For the clarity of exposition, let's restrict this consideration to single-input single-output systems.

A. Problem formulation in MRAC

The conventional problem formulation in MRAC is given as follows:

Determine a state feedback control law $u(t)$ such that $y(t)$ tracks a reference input $r(t)$ **asymptotically** according to the specifications of a given desired system:

$$\begin{aligned} \dot{x}_m(t) &= A_m x_m(t) + b_m r(t)\,, \qquad x_m(0) = x_0 \\ y_m(t) &= c^\top x_m(t)\,, \end{aligned} \tag{2}$$

where $b_m = bk_g$, and $k_g \in \mathbb{R}$ is selected to ensure a unity DC gain of this desired system.

This problem formulation leads to the definition of the **nominal** (*non-implementable*) **controller**:

$$u_{nom}^{MRAC}(t) = \theta^\top(t)x(t) + k_g r(t)\,, \tag{3}$$

the substitution of which into (1) immediately yields (2). The MRAC solution seeks to determine estimation laws for $\theta(t)$, which can be substituted into the nominal controller to yield its adaptive counterpart:

$$u_{MRAC}(t) = \hat{\theta}^\top(t)x(t) + k_g r(t)\,. \tag{4}$$

The determination of $\hat{\theta}(t)$ follows from Lyapunov analysis, which also implies that for constant θ and constant r one has asymptotic convergence $y(t) \to r$, [8]. The estimation rate of $\hat{\theta}(t)$ is referred to as *adaptive gain* in the literature.

B. Problem formulation in $\mathcal{L}_1$ adaptive control

The conventional problem formulation in $\mathcal{L}_1$ adaptive control theory is given as follows:

Determine a state feedback control law $u(t)$ such that $y(t)$ tracks the reference input $r(t)$ and ensures:

- *a priori* **predictable** performance **for all** $t \geq 0$ for different initial conditions and reference signals;
- the control signal $u(t)$ **does not exceed the available bandwidth of the control channel**.

Notice that the first bullet of the problem formulation in $\mathcal{L}_1$ adaptive control asks for specifications from the beginning of the system operation by enforcing *for all* $t \geq 0$, while in MRAC the objective is stated *only asymptotically*. A more important point is the second bullet regarding the bandwidth of the control signal. This second bullet incorporates the robustness specifications for the control system explicitly in the problem formulation of $\mathcal{L}_1$ adaptive control. To understand this requirement let's first look at what can be achieved in the presence of *known* $\theta(t)$ by the nominal controller (3) of MRAC. This controller aims at perfect cancelation of uncertainties; to see this, we note that mathematical substitution of (3) into (1) leads to the desired dynamics (2). In a real application, the unknown parameter vector $\theta(t)$ may have frequencies that lie outside the control channel bandwidth. In this case, these frequencies will be naturally attenuated by the actuator of the system, which implies that –even in the presence of perfect knowledge of $\theta(t)$– the behavior of (2) may not be achievable for any $t \geq 0$.

Based on the understanding that some classes of uncertainties may not be perfectly canceled even if they are completely known, the problem formulation in $\mathcal{L}_1$ adaptive control is *changed*; the new problem formulation in $\mathcal{L}_1$ adaptive control asks for **compensation of uncertainties *only* within the bandwidth of the control channel**, specified by some bandwidth-limited filter $C(s)$; namely, we formulate the **nominal controller in $\mathcal{L}_1$ adaptive control theory** by resorting to the following structure:[1]

$$u_{nom}^{\mathcal{L}_1}(t) = C(s)[u_{nom}^{MRAC}](t) = C(s)[\theta^\top x + k_g r](t)\,, \tag{5}$$

where we use the time-domain notation $y(t) = G(s)[u](t)$ to represent the system $y(s) = G(s)u(s)$. We note that, given a known $\theta(t)$, the filter $C(s)$ can be designed so that the signal in (5) always stays within the control channel bandwidth, irrespective of the frequency content of $\theta(t)$. Obviously, $C(s)$ will attenuate the frequency components in $\theta(t)$ that are above its bandwidth. This definition of nominal controller leads to a different –modified– reference system ($\mathcal{L}_1$ reference system):

$$\begin{aligned} \dot{x}_{ref}(t) &= A_m x_{ref}(t) + b(C(s)[\theta^\top x_{ref} + k_g r](t) - \theta^\top(t)x_{ref}(t))\,, \qquad x_{ref}(0) = x_0, \\ y_{ref}(t) &= c^\top x_{ref}(t)\,, \end{aligned} \tag{6}$$

[1]The reference input can be left outside the filter.

which is *achievable* from $t = 0$. The stability of this reference system can be analyzed by resorting to the $\mathcal{L}_1$-norm condition

$$\|(1 - C(s))(s\mathbb{I} - A_m)^{-1}b\|_{\mathcal{L}_1}\theta_{\max} < 1\,,$$

where $\theta_{\max}$ refers to the maximum value of the $\mathcal{L}_1$ norm of $\theta(t)$.[2]

Given the nominal controller in (5), its corresponding $\mathcal{L}_1$-adaptive counterpart is defined as follows:[3]

$$u_{\mathcal{L}_1}(t) = C(s)[\hat{\theta}^\top x + k_g r](t)\,. \tag{7}$$

For derivation of the adaptive laws, along with the proofs of stability and performance, one can refer to [4, pp. 18–27].

Remark 1: The key point in the above discussion is that the $\mathcal{L}_1$ problem formulation prioritizes the *feasibility of the control objective* with the understanding of the hardware constraints (control channel bandwidth), [9][4]. The MRAC formulation, on the other hand, attempts to compensate for all uncertainties, irrespective of their frequency content.

II. Key Ideas and Assumptions: Role of Architectures

A. Key Idea: Decoupling

The treatment above illustrates the difference in the problem formulations of MRAC and $\mathcal{L}_1$ adaptive control. The key, however, is how **to synthesize an architecture for this new problem formulation that *decouples* the estimation loop from the control loop**. As it will become clear, this decoupling is the critical feature that allows us to increase the adaptive gain in the adaptation laws without hurting the robustness margins (time-delay margin) of the closed-loop adaptive system. It is evident that "the mere introduction of a filter" (as described in [1]) –in series with the adaptive controller– will not achieve this decoupling and will only hurt the robustness properties (time-delay margin) of the closed-loop system. In this section, we describe the basic architecture of $\mathcal{L}_1$ adaptive controllers and emphasize the importance of inserting the **bandwidth-limited filter $C(s)$ in a very particular point of the architecture.**

For this purpose, we consider two adaptive control schemes: a predictor-based MRAC and the corresponding $\mathcal{L}_1$ adaptive controller; see Figure 1. Both architectures use a *state predictor*, which mimics the structure of the system from (1), with the only difference that the unknown parameters are replaced by their corresponding estimates:

$$\begin{aligned} \dot{\hat{x}}(t) &= A_m\hat{x}(t) + b(u(t) - \hat{\theta}^\top(t)x(t))\,, \qquad x(0) = x_0, \\ \hat{y}(t) &= c^\top\hat{x}(t)\,. \end{aligned} \tag{8}$$

It is important to emphasize that a perfect initialization of the predictor, using x_0, is impossible in practice. The non-zero initialization errors are proven to lead to additional exponentially decaying terms in the overall performance bounds, [4]. For brevity, we do not elaborate on these details here. The error signal $\tilde{x}(t) = \hat{x}(t) - x(t)$ is used for synthesis of the adaptive law (estimation of $\hat{\theta}(t)$). Notice that the dynamics of $\tilde{x}(t)$ are **independent of the control law** and are the same for both MRAC and $\mathcal{L}_1$ adaptive control laws:

$$\begin{aligned} \dot{\tilde{x}}(t) &= A_m\tilde{x}(t) - b(\hat{\theta}(t) - \theta(t))^\top x(t)\,, \qquad x(0) = 0, \\ \tilde{y}(t) &= c^\top\tilde{x}(t)\,. \end{aligned} \tag{9}$$

Finally, notice that the closed-loop predictor with the MRAC control law from (4) results in the desired dynamics, given by (2).

It is important to notice from Figure 1b that, in $\mathcal{L}_1$ control architectures, the filtered control signal is sent to both the plant *and* the state predictor; the filter is thus embedded into the adaptive architecture, which implies that **one cannot analyze the closed-loop system as if the adaptive controller and the filter were acting *in series***. The fact that both the predictor and the plant are using the same filtered control signal is crucial to ensure the decoupling of the estimation and control loops. In other words, the presence of the filter at this particular point of the architecture

[2] We notice that the stability of this reference system can be proved using any other norm. However, for the purpose of discussing the transient behavior of the resulting closed-loop system in terms of control specifications, we need to use the $\mathcal{L}_\infty$ norms of signals; hence, since the $\mathcal{L}_1$ norm of a system is the induced $\mathcal{L}_\infty$ norm of its input/output signals, we use the $\mathcal{L}_1$ norm in order to facilitate the transient analysis.

[3] The reference input can be left outside the filter.

[4] "...Note that the available bandwidth is not a function of the compensator or of the control design process. Rather, it is an a priori constraint imposed by the physical hardware we use in the control loop. Most importantly, the available bandwidth is always finite", [9].

is what allows us to increase the adaptive gain without hurting the robustness margins (time-delay margin) of the closed-loop adaptive system. The same filter inserted at a different point (output of the adaptation laws, input of the state predictor only, input of the plant only...) will ***not*** achieve the decoupling of the two loops.

To give more insights into this point, we consider the MRAC architecture in Figure 1a and note that the parameter estimate generated by the estimation loop of MRAC propagates through the control law directly into the plant. Thus, in MRAC architectures, the adaptive gain acts as a feedback gain. Therefore, the speed of adaptation, defined by the adaptation gain (estimation rate), plays a crucial role in the robustness and performance tradeoff in MRAC. Instead, in $\mathcal{L}_1$ adaptive control architectures, the adaptive gain does ***not*** act as a feedback gain, which implies that –*from an architectural perspective*– one can increase the adaptive gain without reducing the robustness margins (time-delay margin) of the closed-loop system. This result is proven in [10]. In short, in $\mathcal{L}_1$ adaptive controllers, the architecture does ***not*** limit the choice of the adaptive gains, whereas in MRAC it does, because high adaptive gains lead to high-gain feedback.

Moreover, we note that the fast estimation loop in $\mathcal{L}_1$ architectures (see Figure 1b) is implemented as a unique block in a processor/computer. The only uncertainties encountered internally are computational and communication delays, cycle time, quantization, etc., and are solely due to the numerical implementation of the controller. It is clear that these 'cyber' uncertainties will put constraints on the choice of the adaptive gain in $\mathcal{L}_1$ control architectures to ensure that the fast estimation loop is numerically stable.

The analysis above illustrates how, for the class of systems in (1) with known input gain, a simple filtering structure –*inserted at the right point within the architecture*– decouples the estimation loop from the control loop, which in its turn allows to increase the adaptive gain arbitrarily, *subject only to hardware limitations*.

B. Assumptions and Extension to Systems with Unknown Input Gain

The next natural question is how far one can move with this approach. In fact, to be honest, the greatest challenge in our past eight years' history was to extend this solution to systems with unknown input gain, subject to the assumption of known sign commonly adopted in the literature, [11]. The reason that I emphasize this point is that, quite often, $\mathcal{L}_1$ adaptive control is viewed by the critics as an input filtered MRAC, as the authors of [1] have attempted to present and interpret without understanding the significance of decoupling.

In general, $\mathcal{L}_1$ adaptive control should ***not*** be perceived just as a result of inserting a bandwidth-limited filter into an MRAC architecture, but rather as a methodology for the synthesis of a **filtering structure that decouples the estimation loop from the control loop**, with analytically provable guarantees. Let's see, for example, why the control law in (7) (presented in Figure 1b) does not work in the presence of unknown input gain.

Consider the dynamics of the same system as in (1) but with unknown input gain ω:

$$\dot{x}(t) = A_m x(t) + b(\omega u(t) + \theta^\top(t) x(t)), \quad x(0) = x_0, \tag{10}$$

where the sign of ω as well as upper and lower bounds on ω are known.

In a predictor-based MRAC scheme, presented in Figure 1a, we estimate the plant parameters $\hat{\theta}(t)$ and $\hat{\omega}(t)$ using the following state predictor:

$$\dot{\hat{x}}(t) = A_m \hat{x}(t) + b\left(\hat{\omega}(t)u(t) + \hat{\theta}^\top(t)x(t)\right), \qquad \hat{x}(0) = x_0, \tag{11}$$

which mimics the structure of the plant (10). Using the same *certainty equivalence principle* that we used to derive (4), the MRAC control law is given by:

$$u(t) = \frac{1}{\hat{\omega}(t)}\left(-\hat{\theta}^\top(t)x(t) + k_g r(t)\right). \tag{12}$$

Substituting (12) into (11), we obtain the desired system (2). Notice that this definition of the control law requires the adaptation law to ensure that the estimate $\hat{\omega}(t)$ remains bounded away from zero. For this purpose one can use projection-based adaptation laws [12].

In this case, however, the $\mathcal{L}_1$ adaptive control law cannot be defined just as the filtered version of the MRAC control law similar to what we did in the previous section. To show where the problem lies, consider the filtered version of the control signal in (12), similar to (7):

$$u_f(s) = C_f(s)u(s), \tag{13}$$

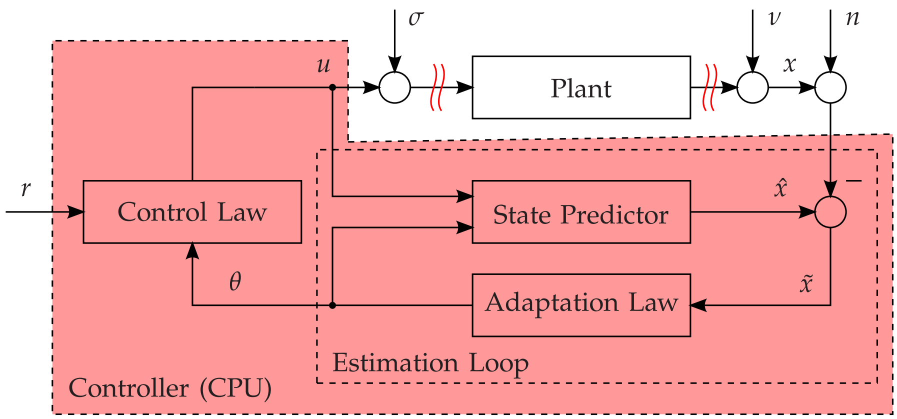

(a) MRAC

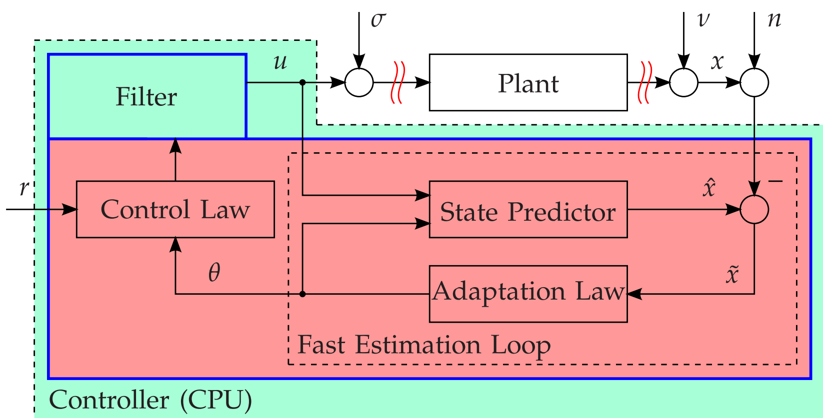

(b) $\mathcal{L}_1$ adaptive controller

Fig. 1: Adaptive control architectures. The state predictor mimics the system structure with the unknown parameters replaced by the parameter estimates. The parameter estimates generated by the estimation loop of MRAC propagate through the control law directly into the plant. In the presence of fast estimation rates, the control signal may exceed the available control system bandwidth, which can hurt the robustness of the closed-loop adaptive system. Therefore, the speed of adaptation, defined by the estimation rate, plays a crucial role in the robustness and performance tradeoff for MRAC schemes. In $\mathcal{L}_1$ adaptive controllers, *the estimation loop is decoupled from the control loop*; the filter protects the control signal from exceeding the available control system bandwidth, as it shields the high-frequency content of the parameter estimates, due to the large rates, inside the controller block. The fast estimation loop of $\mathcal{L}_1$ adaptive controller is implemented completely in an *uncertainty-free environment* and, therefore, it allows fast estimation rates without hurting robustness. We also notice that the loop breaking points used in conventional robust control cannot be affected by the large gains of the fast estimation loop of $\mathcal{L}_1$ adaptive controller.

where $u(s)$ is the Laplace transform of $u(t)$ in (12), and $C_f(s)$ is a low-pass filter. Let $c_f(t)$ be the impulse response of the transfer function $C_f(s)$. Then

$$u_f(t) = c_f(t) * u(t) = c_f(t) * \left(\frac{k_g r(t) - \hat{\theta}^\top(t) x(t)}{\hat{\omega}(t)} \right),$$

where $*$ denotes the convolution operator[5]. Substituting this expression into the system dynamics (10), we get

$$\dot{x}(t) = A_m x(t) + b \left(\omega \left(c_f(t) * \left(\frac{k_g r(t) - \hat{\theta}^\top(t) x(t)}{\hat{\omega}(t)} \right) \right) + \theta^\top(t) x(t) \right), \qquad x(0) = x_0 , \tag{14}$$

while substituting its ideal (non-adaptive) version into the corresponding $\mathcal{L}_1$ reference system (following the same

[5]The convolution operator $c(t) = a(t)*b(t)$ for signals $a(t)$, $b(t)$, $c(t) \in \mathbb{R}$ for $t \geq 0$ is defined as $c(t) = \int_0^t a(t-\tau)b(\tau)d\tau$.

philosophy as the derivation of (6)) leads to

$$\dot{x}_{ref}(t) = A_m x_{ref}(t) + b\left(\omega\, C_f(s)\left[\frac{k_g r - \theta^\top x_{ref}}{\omega}\right](t) + \theta^\top(t) x_{ref}(t)\right), \qquad x_{ref}(0) = x_0\,. \tag{15}$$

In (15) $\frac{1}{\omega}$ can be moved in front of $C_f(s)$ to cancel ω; however, the convolution operator does ***not*** allow moving $\frac{1}{\hat{\omega}(t)}$ outside the convolution operator in (14) to write $\frac{\omega}{\hat{\omega}(t)}$. This fact renders the derivation of the performance bounds between $x(t)$ and $x_{ref}(t)$ (and also $u(t)$ and $u_{ref}(t)$) overly complicated, resulting in a cumbersome and inelegant proof (to say the least). As a consequence, predictability of the control law (13) is questionable.

The $\mathcal{L}_1$ adaptive controller for this system is defined using the following filtering structure (see Figure 2):

$$u(s) = \frac{k}{s}\hat{\eta}(s)\,, \tag{16}$$

where $\hat{\eta}(s)$ is the Laplace transform of

$$\hat{\eta}(t) = k_g r(t) - \hat{\theta}^\top(t) x(t) - \hat{\omega}(t) u(t)\,.$$

The control law in (16) can thus be expressed as:

$$\dot{u}(t) = -k\hat{\omega}(t)u(t) + k\hat{\omega}(t)\left(\frac{k_g r(t) - \hat{\theta}^\top(t)x(t)}{\hat{\omega}(t)}\right), \qquad u(0) = 0\,, \tag{17}$$

and, therefore, $\hat{\omega}(t)$ can be viewed as a time-varying gain affecting the 'bandwidth' of a first-order linear time-varying filter. This approach (even if it might not seem trivial from this description) allows to derive a rather simple proof of the performance bounds of this adaptive architecture. For details, see [4]. If we consider the case of slowly varying reference signals, equation (17) gives us further insights into this control law. In steady state, when $\dot{u}(t) \approx 0$, we have that

$$k_g r(t) - \hat{\theta}^\top(t) x(t) - \hat{\omega}(t) u(t) \approx 0\,.$$

Comparing the control law in (16) to the indirect MRAC control law (12), we find that the $\mathcal{L}_1$ control law avoids division by $\hat{\omega}(t)$. The filter in the $\mathcal{L}_1$ control law (16) solves the design equation dynamically by driving $\hat{\eta}(t)$ to zero.[6]

Remark 2: Notice that in the definition of the filtered control law (16), in general, one can use an arbitrary strictly-proper transfer function $D(s)$, which leads to a higher-order, strictly-proper, stable filter with extra design parameters:

$$C(s) = \frac{k\omega D(s)}{1 + k\omega D(s)}\,. \tag{18}$$

Also, if $D(s)$ includes an integrator, then $C(s)$ has a unity DC gain, i.e. $C(0) = 1$, [4]. In this case, the control law in (16) takes the following form [4]:

$$u(s) = kD(s)\hat{\eta}(s)\,. \tag{19}$$

C. Simulation Example

To illustrate the decoupling of the estimation loop from the control loop for the filtering structure presented above, we refer to the simulation example from [1]. Simulation files are available upon request.

1) $\mathcal{L}_1$ Adaptive Controller for the Nominal System: Consider the following scalar system, [1]:

$$\dot{x}(t) = -x(t) + \omega u(t)\,, \qquad x(0) = x_0\,, \tag{20}$$

where ω is an unknown parameter with known sign and lower and upper bounds. The $\mathcal{L}_1$ adaptive controller for this system consists of the state predictor

$$\dot{\hat{x}}(t) = -\hat{x}(t) + \hat{\omega}(t)u(t)\,, \qquad \hat{x}(0) = x_0\,,$$

where the estimate $\hat{\omega}(t)$ is governed by the following adaptation law:

$$\dot{\hat{\omega}}(t) = \Gamma\,\mathrm{Proj}(\hat{\omega}(t), -\tilde{x}(t)u(t))\,, \qquad \hat{\omega}(0) = \hat{\omega}_0\,,$$

[6]This insight was shared with us by Karl J. Åström.

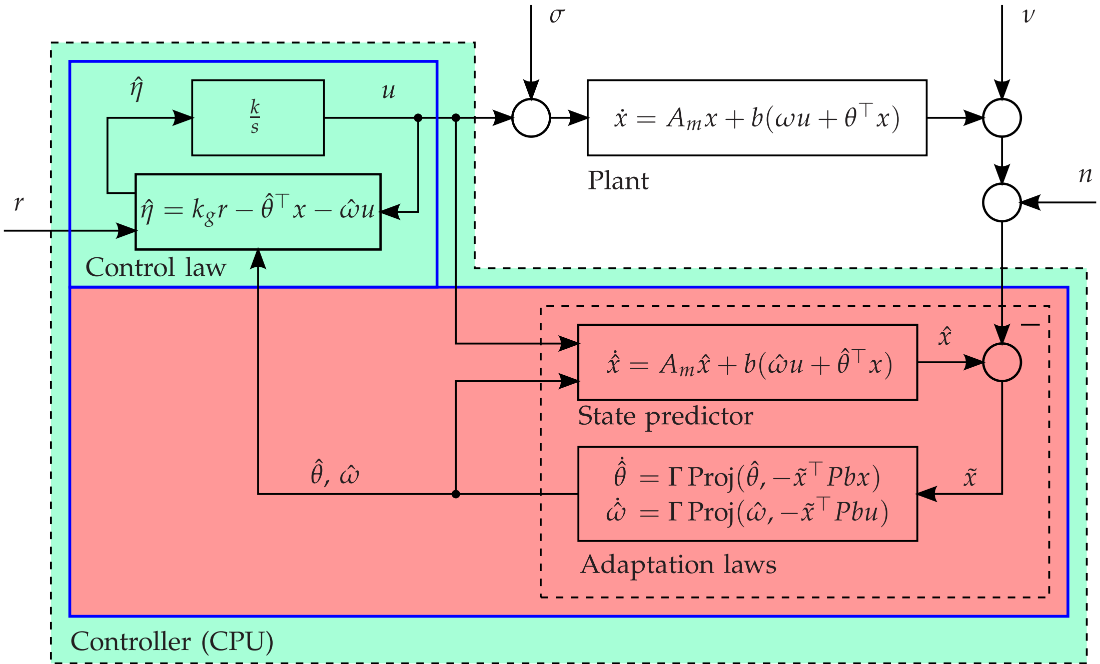

Fig. 2: $\mathcal{L}_1$ adaptive control architecture for systems with unknown input gain. The estimation loop is embedded inside the controller and is fully decoupled from the system uncertainties. The filtering structure for synthesizing the feedback law is shielding the high frequencies due to fast estimation rates from propagating into the system.

and the control signal is generated as

$$u(s) = kD(s)\hat{\eta}(s)\,,$$
$$\hat{\eta}(t) = r(t) - \hat{\omega}(t)u(t)\,.$$

Simulation results are shown in Figure 3 for $\omega = 1$ (as in [1]), $x_0 = 0$, and $\hat{\omega}_0 = 0.5$. For the design of the $\mathcal{L}_1$ adaptive controller we choose $k = 5$ and $D(s) = \frac{1}{s(0.05s+1)}$. The adaptation gain is set to $\Gamma = 500{,}000$, and the projection bounds are set to $\hat{\omega}(t) \in [0.1,\ 2\Delta]$, $\hat{\theta}(t) \in [-\Delta,\ \Delta]$, $\hat{\sigma}(t) \in [-\Delta, \Delta]$, with $\Delta = 10^6$. As can be observed in Figure 3, the system output tracks the reference command (with zero steady-state error) according to the desired specifications. The control signal is smooth and does not have any spikes or oscillations, despite the high frequencies in the estimated parameters due to the large estimation gain (fast adaptation).

2) $\mathcal{L}_1$ Adaptive Controller in the Presence of Unmodeled Dynamics: Next, we check robustness of this $\mathcal{L}_1$ adaptive controller for the unmodeled input dynamics suggested in [1]. For this purpose, we rewrite (20) as follows:

$$\dot{x}(t) = -x(t) + \omega\mu(t)\,, \qquad x(0) = x_0\,,$$

where $\mu(s) = (1 - 0.02s)u(t)$ is the output of the unmodeled dynamics at the system's input. We use the same $\mathcal{L}_1$ adaptive controller as above, and do not do any redesign or retuning. The simulation results in Figure 4 show that in the presence of the unmodeled dynamics the system performance is very close to the original performance without the unmodeled dynamics, given in Figure 3. We also observe that, after the initial transient, the parameter estimate remains in the neighborhood of the actual value of the system input gain and does not hit the projection bounds. The fact that the unmodeled dynamics only lead to a minor degradation of performance and do not *destabilize* the system clearly shows that the large adaptive gain does not lead to high-gain feedback control, implying in its turn that the filtering structure of $\mathcal{L}_1$ adaptive control effectively decouples the estimation loop from the control loop.

3) $\mathcal{L}_1$ Adaptive Controller with Three Adaptive Parameters: We notice that, while the $\mathcal{L}_1$ adaptive controller shown above is suitable for the system in (20) and exhibits good robustness properties, it may not be practical if the system has additional uncertainties besides the input unmodeled dynamics proposed by the authors of [1]. Next we present a more general $\mathcal{L}_1$ adaptive control structure from [4, Section 2.3], which adapts for three types of

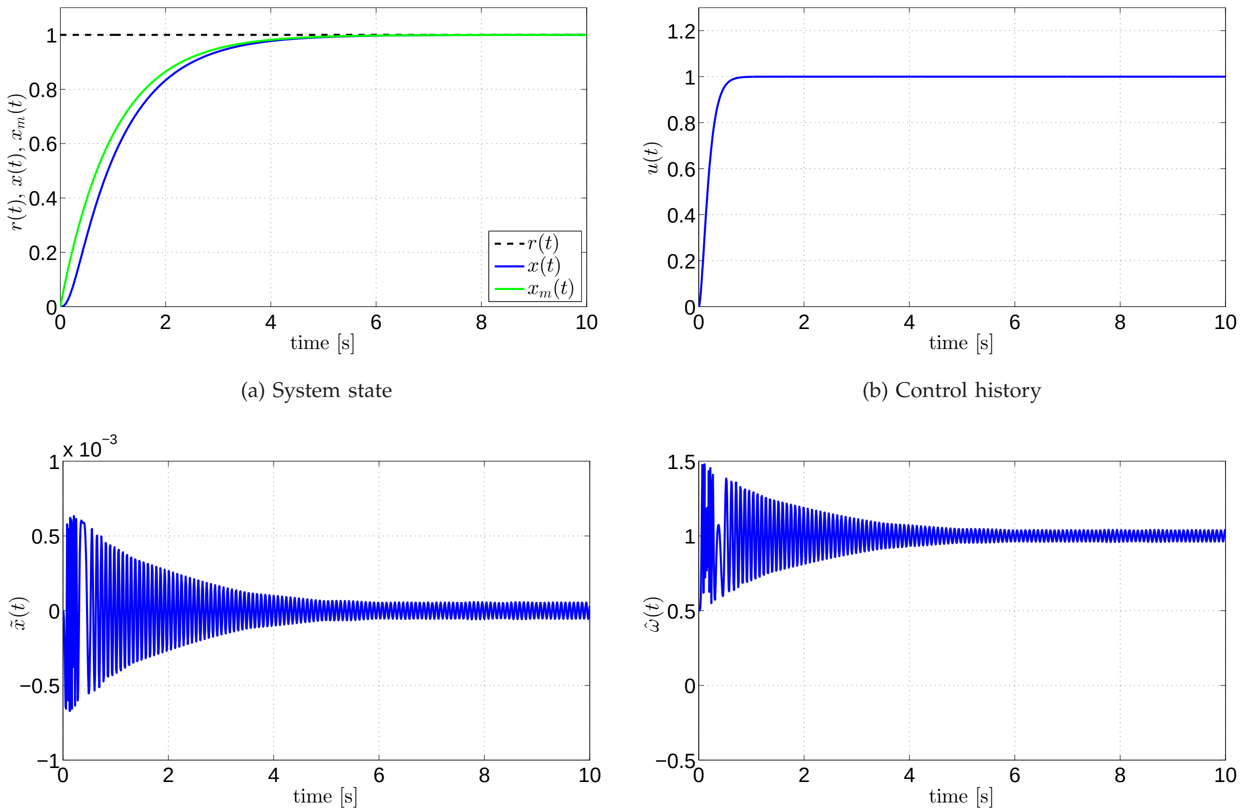

(a) System state

(b) Control history

(c) Prediction error

(d) Parameter estimates

Fig. 3: Simulation results for the $\mathcal{L}_1$ adaptive controller without unmodeled dynamics. The closed-loop system output tracks the reference command according to the desired specifications. The control signal is smooth and does not have any spikes or oscillations. The parameter estimate experiences high-frequency oscillations, which do not propagate into the control signal because of the decoupling estimation and control loops in the $\mathcal{L}_1$ adaptive control architecture.

parametric uncertainties. It consists of the following state predictor:

$$\dot{\hat{x}}(t) = -\hat{x}(t) + \hat{\omega}(t)u(t) + \hat{\theta}(t)x(t) + \hat{\sigma}(t)\,, \qquad \hat{x}(0) = x_0\,, \tag{21}$$

where the parameter estimates are generated using the following projection-based adaptation laws:

$$\dot{\hat{\omega}}(t) = \Gamma\,\mathrm{Proj}(\hat{\omega}(t), -\tilde{x}(t)u(t))\,, \qquad \hat{\omega}(0) = \hat{\omega}_0\,, \tag{22}$$

$$\dot{\hat{\theta}}(t) = \Gamma\,\mathrm{Proj}(\hat{\theta}(t), -\tilde{x}(t)x(t))\,, \qquad \hat{\theta}(0) = \hat{\theta}_0\,, \tag{23}$$

$$\dot{\hat{\sigma}}(t) = \Gamma\,\mathrm{Proj}(\hat{\sigma}(t), -\tilde{x}(t))\,, \qquad \hat{\sigma}(0) = \hat{\sigma}_0\,, \tag{24}$$

and the $\mathcal{L}_1$ adaptive control signal is generated as

$$u(s) = kD(s)\hat{\eta}(s)\,,$$
$$\hat{\eta}(t) = r(t) - \hat{\sigma}(t) - \hat{\theta}(t)x(t) - \hat{\omega}(t)u(t)\,.$$

Simulation results for the system in the presence of unmodeled input dynamics are shown in Figure 5, where we use $\omega = 1$, $x_0 = 0$, $\hat{\omega}_0 = 0.5$, $\hat{\theta}_0 = \hat{\sigma}_0 = 0$. For the design of the $\mathcal{L}_1$ adaptive controller, we use the same control parameters as in the previous scenario. We see that the control system achieves tracking of the step reference command (with zero steady-state error) according to the desired specifications. The control signal is smooth and does not experience oscillations. The prediction error is relatively small for all time $t \geq 0$, and the parameter estimates remain bounded without hitting the projection bounds.

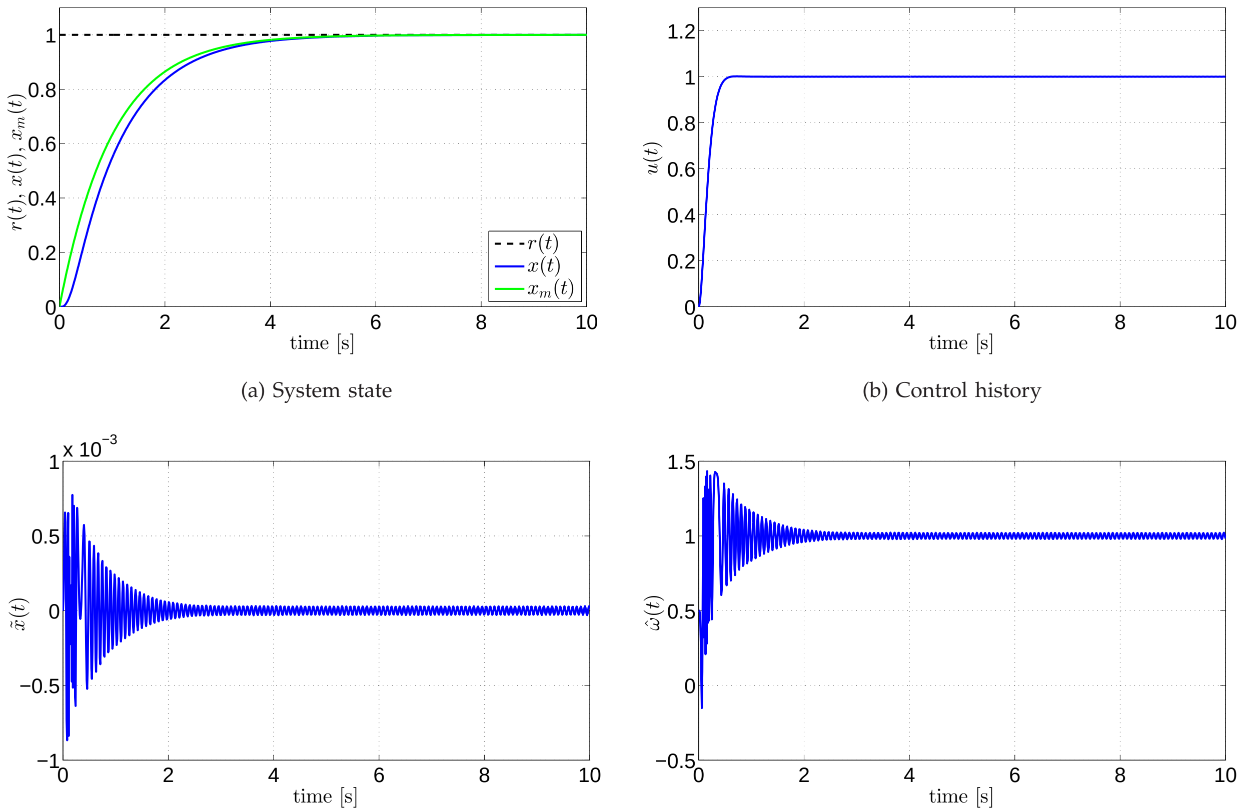

(a) System state

(b) Control history

(c) Prediction error

(d) Parameter estimates

Fig. 4: Simulation results for the $\mathcal{L}_1$ adaptive controller in the presence of unmodeled input dynamics. Both the system output and the control signal remain close to the signals obtained in the absence of the unmodeled dynamics in Figure 3. The parameter estimate remains in the neighborhood of the actual value and does not hit the projection bound.

4) Comparison of the Time-delay Margins of $\mathcal{L}_1$ Adaptive Controller and MRAC: Finally, we test the robustness of the $\mathcal{L}_1$ adaptive controller to time delays and compare it to the corresponding MRAC architecture. For this purpose, we consider the system (20) and the $\mathcal{L}_1$ adaptive controller with three parameter estimates. We use the same control parameters as above for the design of the $\mathcal{L}_1$ adaptive controller.

For the design of MRAC we use the same state predictor and adaptation laws as in (21)–(24). The control law of MRAC is given by

$$u(t) = \frac{1}{\hat{\omega}(t)} \left(r(t) - \hat{\theta}(t)x(t) - \hat{\sigma}(t) \right) .$$

Next we compute numerically the time-delay margin for both controllers as a function of the adaptation gain. The results in Figure 6 show that the time-delay margin of MRAC vanishes as the adaptation gain is increased, while the time-delay margin of $\mathcal{L}_1$ adaptive controller remains bounded away from zero. For large adaptation gains it approaches the value $\mathcal{T} = 0.28$ s.

Next we show the transient behavior of the $\mathcal{L}_1$ adaptive controller in the presence of time delays. For this purpose, we consider a time delay at the system input of $\tau = 0.1$ s (comparable to the time-delay margin of the system). The simulation results are given in Figure 7. As we see, the transient response of the closed-loop system with the $\mathcal{L}_1$ adaptive controller is very close to the response without time-delay shown in Figure 5.

5) Scaled Transient Response of $\mathcal{L}_1$ Adaptive Controller: Next, we compare the transient responses of $\mathcal{L}_1$ adaptive controller to MRAC. For this purpose, we use the same MRAC as above and set its adaptation gain to $\Gamma_{MRAC} = 0.9$, which results in the same time-delay margin as it is achieved by the $\mathcal{L}_1$ adaptive controller in the presence of fast adaptation. The simulation results for 3 step commands of different amplitude in the presence of unmodeled

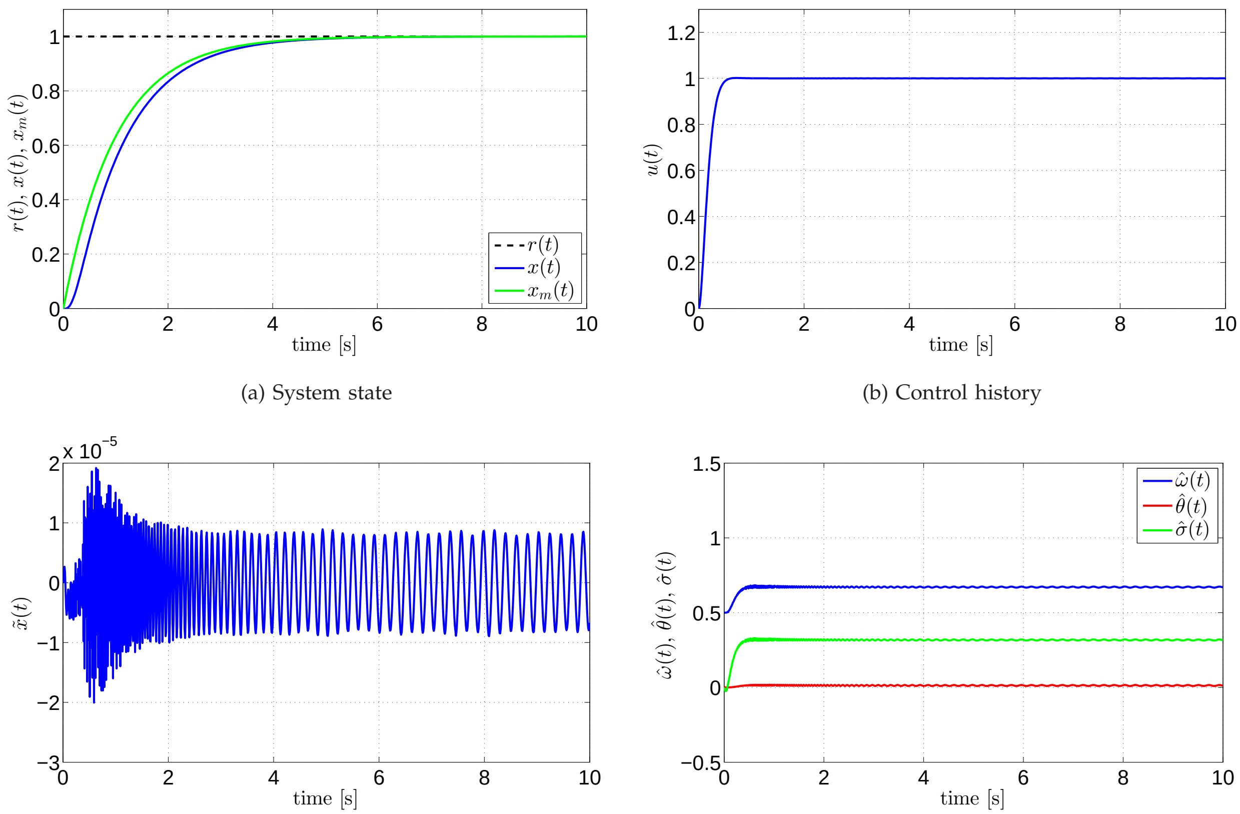

(a) System state

(b) Control history

(c) Prediction error

(d) Parameter estimates

Fig. 5: Simulation results for the $\mathcal{L}_1$ adaptive controller with projection based adaptation laws. The system state achieves tracking of the step reference command according to the desired specifications in the presence of unmodeled input dynamics. The control signal is smooth and does not experience oscillations.

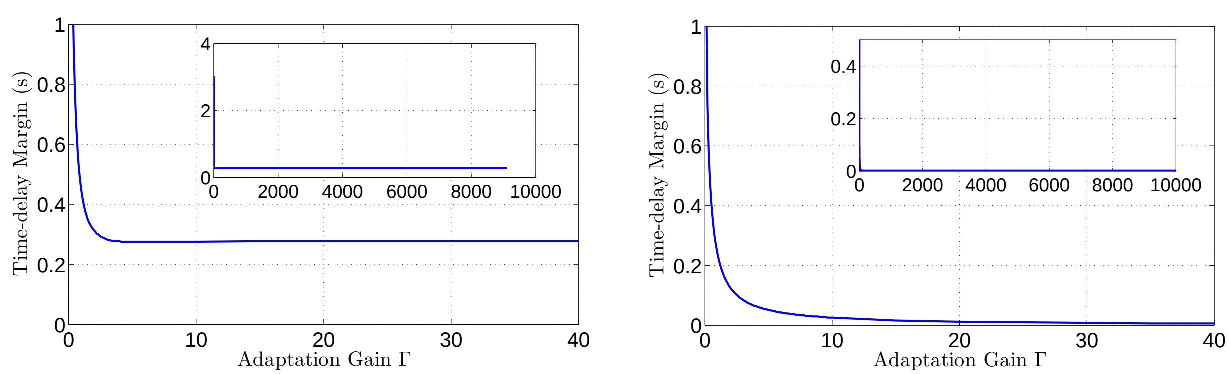

(a) $\mathcal{L}_1$ adaptive controller

(b) MRAC

Fig. 6: Time-delay margin of $\mathcal{L}_1$ adaptive controller and MRAC as a function of adaptation gain. The time-delay margin of MRAC vanishes as the adaptation gain is increased, while the time-delay margin of $\mathcal{L}_1$ adaptive controller remain bounded away from zero as the adaptation gain increases. For large adaptation gains it approaches the value $\mathcal{T} = 0.28$ s.

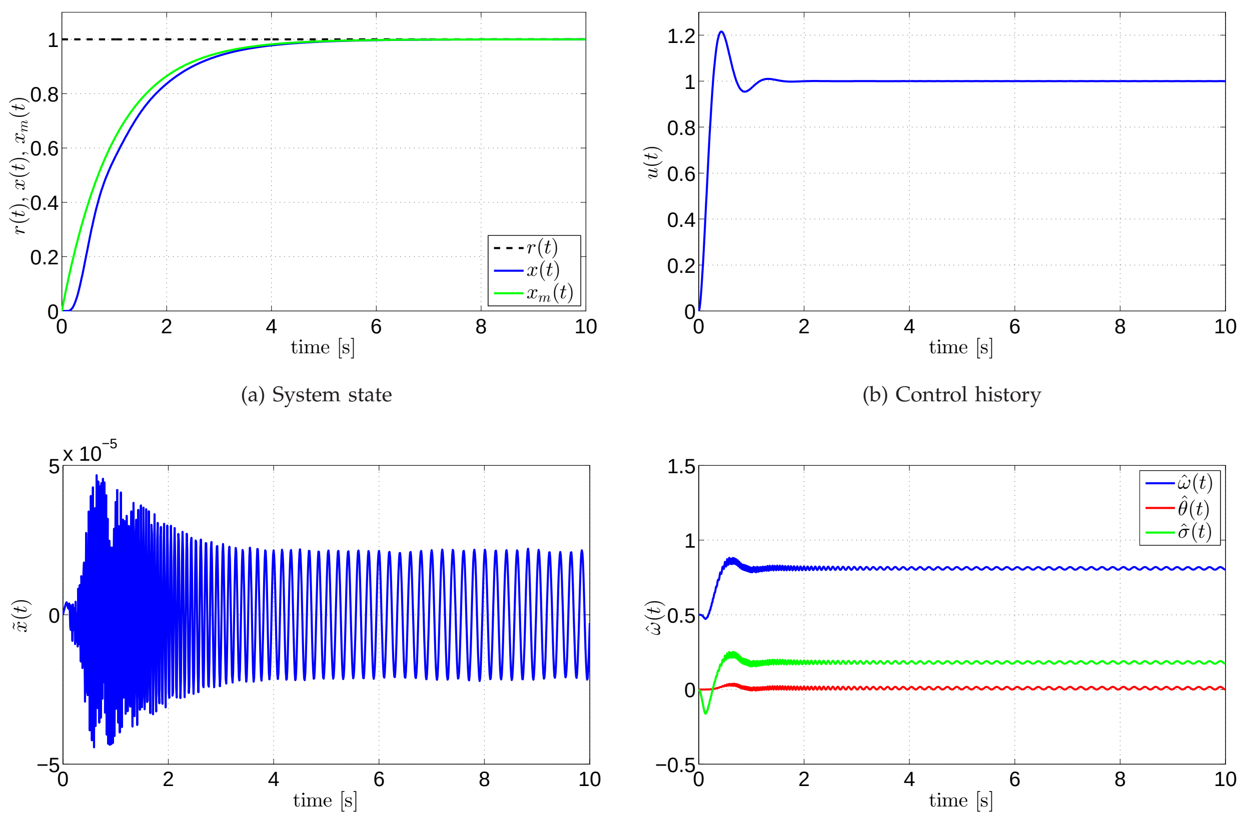

(a) System state

(b) Control history

(c) Prediction error

(d) Parameter estimates

Fig. 7: Simulation results for the $\mathcal{L}_1$ adaptive controller in the presence of time delay $\tau = 0.1$ s (the time delay margin of the control system is approximately 0.28 s). The system input and output performance are very close to the results shown in Figure 5.

input dynamics are given in Figure 8. We see that the state response of the $\mathcal{L}_1$ adaptive controller scales with the amplitude of the step command (similar to linear systems), whereas the transient response of MRAC changes in an unpredictable manner.

Remark 3: With regards to authors' claim in [1] regarding the roles of the projection operator and the filter in the stability and performance of $\mathcal{L}_1$ adaptive controller, we would like to mention that in $\mathcal{L}_1$ adaptive control design and analysis both are equally important. In particular, the above example demonstrates that with the projection operator inactive (by selecting sufficiently large projection bounds), the system is robust to the unmodeled dynamics.

D. $\mathcal{L}_1$ Adaptive Architectures for Various Classes of Systems in State Feedback and Output Feedback

A natural question arises as how to synthesize the filtering structure for various classes of systems, including strict-feedback systems with nonlinearities, systems with unmodeled dynamics, unmatched uncertainties, using state feedback and output feedback. For example, how to handle non-SPR systems in an output feedback setting? A systematic way for developing all these architectures and the corresponding filtering structures, including special adaptive laws for non-SPR systems with output feedback, can be found in [4].

Remark 4: Notice that there is no universal $\mathcal{L}_1$ adaptive controller that would fit all the systems, as there is no universal MRAC controller that would fit all the systems. Synthesizing a correct filtering structure that *decouples the estimation loop from the control loop* is the critical element to understand the development of the $\mathcal{L}_1$ adaptive control theory. As an example, we refer to [13], where the author naively extended the input filtering from Section I, according to Figure 1b, to Monopoli's architecture from [14] to conclude that the $\mathcal{L}_1$ adaptive controller cannot achieve perfect tracking in the presence of bounded disturbances. In response, we published [15], where we showed

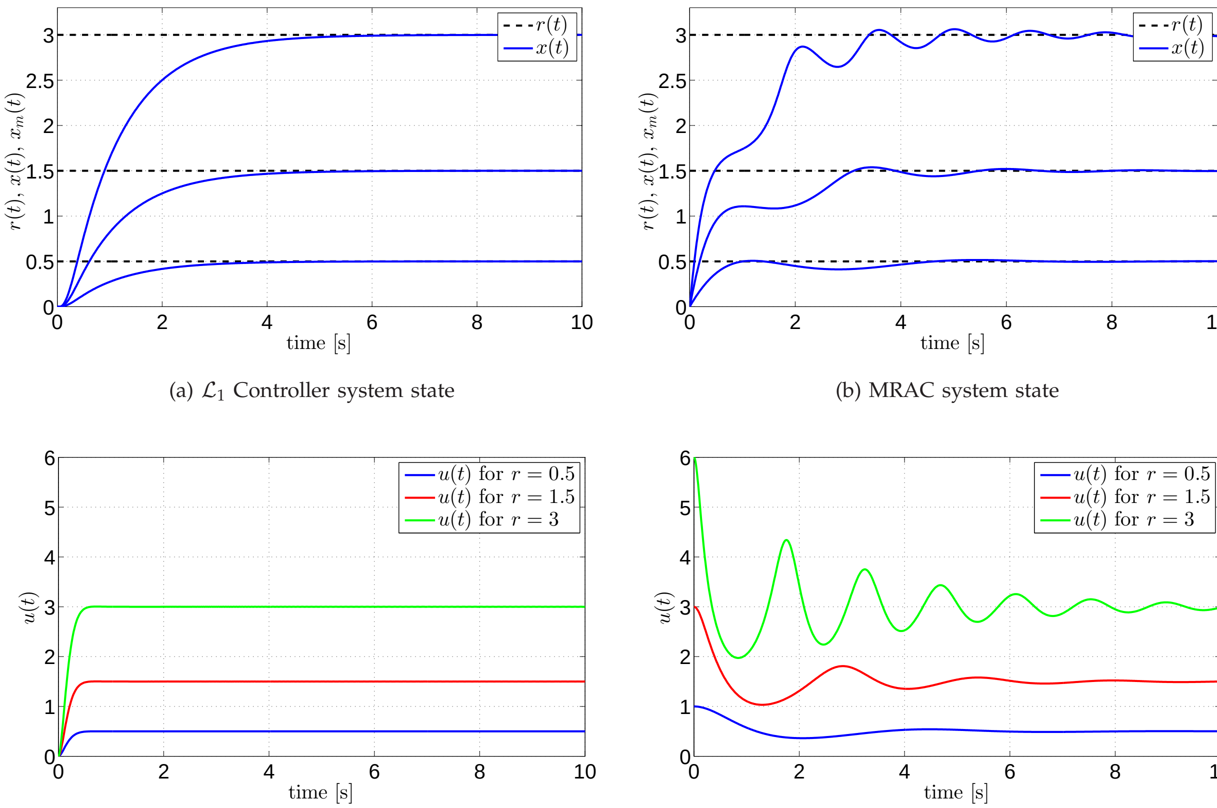

(a) $\mathcal{L}_1$ Controller system state

(b) MRAC system state

(c) $\mathcal{L}_1$ Controller control signal

(d) MRAC control signal

Fig. 8: Comparison of transient responses of $\mathcal{L}_1$ adaptive controller and MRAC in the presence of unmodeled dynamics. The state response of the $\mathcal{L}_1$ adaptive controller scales with the magnitude of the step command, whereas the transient response of MRAC is highly nonlinear and changes in an unpredictable manner.

how to synthesize the filtering structure and the corresponding $\mathcal{L}_1$ architecture, appropriate for that particular Monopoli-style MRAC scheme, to achieve perfect decoupling between the estimation and the control loops.

III. $\mathcal{L}_1$ Adaptive Control Design Challenges

As explained in the preceding sections of this document, the main challenge in synthesizing an $\mathcal{L}_1$ controller for a given plant is to understand how to build the filtering structure so that the estimation loop and the control loop are decoupled. Another open problem is how to choose the parameters of the filtering structure so that the sufficient conditions of stability are verified and a given performance index is optimized; in particular, this design problem is extremely challenging in output feedback architectures. Some partial solutions for addressing these problems are summarized in [16]–[18].

IV. Wrong Claims in "$\mathcal{L}_1$-adaptive control: Stability, robustness, and misperceptions" [1]

In the light of the above discussion, this section rebuts claims and statements in [1]. We analyze the claims as they appear in [1], indicating the section in which a particular claim was made.

Abstract:

- The claim *"… arbitrarily fast convergence to desired values …"* needs clarification. It appears that the authors of [1] have misinterpreted our results in stating "arbitrarily fast convergence to desired values" as one of the properties of our controllers. Their statement suggests parameter convergence, which we never claimed. Instead, in all of our work, we have analyzed the transient and steady-state responses of the input and output signals of the closed-loop adaptive system, and have derived bounds on the maximum deviation of these signals from the corresponding desired response. To be more specific, our claims refer to uniform performance bounds for both the input and the state of the closed-loop system with respect to the corresponding signals of the $\mathcal{L}_1$ reference system, introduced earlier in this document. This $\mathcal{L}_1$ reference system assumes partial compensation of uncertainties within the bandwidth of filter $C(s)$, and hence does not exhibit the ideal performance. This reference system in its turn can approximate the response of the desired (ideal) system by appropriate tuning of the filter, which characterizes the tradeoff between performance and robustness. As explained in detail in [4], the role of this *auxiliary* $\mathcal{L}_1$ reference system is to decouple the analysis between the estimation process and the performance/robustness properties of the closed-loop system.

Section I. Introduction:

- *"Reference [3] was the first instance of an $\mathcal{L}_1$-formulation presented in the literature (see [3], Section 9.4) even though no name was given to the modified scheme"*. We note that Section 9.4 of reference [3] in [1] has, in fact, a formulation that uses the $\mathcal{L}_1$ norm in its analysis; however, the control architecture presented in that section does not have any resemblance to the architectures of the $\mathcal{L}_1$ adaptive control theory. For example, the architecture that the authors present in Section 9.4.1 in' [1] is a combination of a conventional adaptive control scheme with a non-adaptive model-reference PI controller, in which $1/\tau$ acts as both the proportional and the integral gains. The authors argue that tracking performance can be improved "by choosing an arbitrarily small τ". However, the authors fail to mention that, as one reduces the control parameter τ, the non-adaptive PI controller becomes a high-gain feedback control scheme that dominates the adaptive counterpart. Reducing τ results thus in improved performance at the price, however, of reduced robustness margins. The same discussion applies to the controller proposed in Section 9.4.2 in [1], which is just a generalization of the one in Section 9.4.1.
- *"According to the authors of L1-AC, this is accomplished by the mere introduction of an input filter and boosting the adaptive gain of the estimated parameters to very large values"*. This statement is not precise at all. As explained in Section II of this document, the key in the design of $\mathcal{L}_1$ adaptive controllers is the combination of a state predictor, a filtering structure (**inserted at a very precise point of the control architecture**), and the structure of the estimation laws with appropriate bounding features.
- The claims 1) through 4) on page 3 of [1] are ***not*** correct. These four points will be discussed below.

Section II. What is $\mathcal{L}_1$-AC?:

Section II falls short of making complete claims and only presents the most basic $\mathcal{L}_1$ adaptive controller for the most simple class of systems. In particular,

- Regarding 1), one must add that the difference between $\mathcal{L}_1$ and MRAC, represented by the addition of the filter $C(s)$, is true *only* for this simple class of systems. As explained earlier in this document, for other classes of systems, the filtering structure of $\mathcal{L}_1$ adaptive controllers may take a different form, making these controllers significantly different from the existing MRAC architectures.
- Regarding 2), the global stability results obtained with MRAC can also be claimed for $\mathcal{L}_1$ adaptive control.
- Regarding 3), the authors fail to mention that, in addition to the bound on the state tracking error, $\mathcal{L}_1$ adaptive control architectures have a similar bound on the adaptive control signal, whereas MRAC cannot claim this second, *very important*, bound. In particular, the bound on the control signal is critical to analyze the robustness properties of $\mathcal{L}_1$ adaptive controllers. The lack of this uniform bound on the control signal of MRAC leads consequently to lack of robustness in the presence of fast adaptation.

Section III. $\mathcal{L}_1$-AC: Stability:

- The first bullet on page 5 in [1] hides a subtle, but *critical*, difference between the two architectures that the authors fail to mention. Large adaptive gains in MRAC architectures lead to high-gain feedback with reduced robustness margins, which implies that robustness of MRAC architectures is lost as the adaptive gain is increased; instead, in $\mathcal{L}_1$ adaptive control architectures, robustness is decoupled from the adaptive gain (see [4], pp. 8–15 and 47–59). The authors again omit this important point.
- The authors also omit the important bound on the control signal, which MRAC does ***not*** have.
- A similar attitude is taken in the second bullet with respect to the bound in equation (10) of [1]. The authors fail to mention that $\mathcal{L}_1$ adaptive control has a similar bound for the control signal, which is critically important for robustness analysis. Moreover, notice that from equation (10) in [1] it follows that a non-zero initialization error leads to exponentially decaying terms, which is another proof of guaranteed transient response. This result is further articulated in the time-delay margin proof, which, unfortunately, the authors did not follow-up (see [4], pp. 47–59).

- The comment in the last bullet regarding the non-zero steady-state tracking error is not correct, as $\mathcal{L}_1$ adaptive control, similar to MRAC, guarantees zero steady-state tracking error for *constant reference inputs*. Remark 2.1.2 in [4] has the proof of this result, which the authors omit. If the reference input is time-varying, then neither MRAC nor $\mathcal{L}_1$ adaptive control will ensure asymptotic tracking of $r(t)$. However, we notice that, for time-varying reference inputs, $\mathcal{L}_1$ adaptive control may exhibit a larger lag due to the filter than the MRAC controller. Therefore, the design of the filter in $\mathcal{L}_1$ adaptive control has to be done carefully with consideration of *both transient and steady-state* specifications, along with the *robustness requirements*.

Section IV. $\mathcal{L}_1$-AC: Robustness:

- For the system in equation (16) of [1], the $\mathcal{L}_1$ adaptive control is ***not*** given by equation (17) of [1]. This controller proposed by the authors in [1] seems to be a conventional MRAC in series with a low-pass filter. As explained in the earlier sections of this document, this is ***not*** an $\mathcal{L}_1$ adaptive controller. The correct $\mathcal{L}_1$ adaptive control architecture for this system has been presented in Section II-C of this document. Hence, all the results and claims made in Section IV.B of [1] are **plainly wrong**.

Section V. $\mathcal{L}_1$-AC: Output Feedback:

- The authors argue that for the considered class of systems the $\mathcal{L}_1$ controller is a simple integral controller with anti-windup protection. While for the first-order reference systems the structure of the $\mathcal{L}_1$ adaptive output-feedback controller can be interpreted as filtered PI controller with anti-windup, we would like to note a couple of points:
  - First, if we look into the system dynamics in (1) and replace $\theta^\top(t)x(t)$ by $\theta(t)$, then the MRAC controller will *reduce* to a model-following PI controller. Hence, it should not be surprising that in certain cases the $\mathcal{L}_1$ adaptive controller can be also reduced to a simpler structure; and
  - Second, there are important *architectural differences* between a conventional model-following filtered PI and the corresponding $\mathcal{L}_1$ controller. As explained in the beginning of this document (Sections I and II), the architectures of $\mathcal{L}_1$ adaptive control theory cannot be viewed as a filtered MRAC scheme, because the filter is ***not*** cascaded with the MRAC control law. The authors of [1] should ***not*** underestimate the specific location of the filtering structure in the architecture, and how it should be synthesized to effectively *decouple* the estimation and the control loops so that the integral (adaptation) gain can be increased without leading to high-gain feedback.

  Moreover, we notice that, for more general classes of systems, the $\mathcal{L}_1$ adaptive controller remains nonlinear [15]. Hence, the case of $\mathcal{L}_1$ adaptive output-feedback architecture considered in [1] is ***not*** representative. Therefore, the claims made by authors in Section V of [1] cannot be generalized to any other $\mathcal{L}_1$ adaptive output-feedback control architecture.
- The authors refer to reference [27] in [1] to claim that *"simulation studies of L1-AC to flight platforms ... have observed mediocre performance compared to other adaptive algorithms"*. However, it is easy to verify that reference [27] from [1] has a wrong implementation of $\mathcal{L}_1$ adaptive control. This issue was already discussed with the authors of [27], and it is therefore surprising to see that this paper is being used in [1] as a 'failure' of an $\mathcal{L}_1$ adaptive controller.
- Regarding the implementations in [25] and [26] of [1], we agree that the implemented $\mathcal{L}_1$ adaptive controllers reduce to a linear time-varying structure. For the class of systems in [26], we were able to prove (with some delay) that the common nonlinear gradient-minimization-based adaptive laws with projection can also be applied with provable guarantees; this result was included in [4, Section 3.2]. As explained in Remark 3.3.2 of [4], the performance bounds for both adaptive laws appeared to be *identical*, and hence there is no benefit in trying to implement also the gradient-minimization-based adaptive laws with the projection operator. There are, however, several other implementations of $\mathcal{L}_1$ adaptive control laws where the commonly adopted gradient-minimization-based nonlinear adaptive laws have been employed (see, for example, [19]–[21]).

Section VI. Conclusions:

- As clarified earlier in this same section, the control architectures from [3, Section 9.4] in [1], the analysis of which employs the $\mathcal{L}_1$ norm, are ***different*** from the architectures of the $\mathcal{L}_1$ adaptive control theory in [4].
- Asymptotic tracking with zero steady-state error of step inputs is proved in Remark 2.1.2 of [4].

- The presence of the filter embedded in the $\mathcal{L}_1$ adaptive control architecture protects the stability margins and is not affecting adversely the robustness in the presence of large adaptive gain (it might be a good idea to study the proof of Theorem 2.2.4 from [4]).
- We agree that parameter projection is a critical element in the proofs of $\mathcal{L}_1$ adaptive control theory; however, as discussed in the first sections of this document, the presence of the bandwidth-limited filter at a very particular point of the control architecture is what ensures the decoupling between the estimation and the control loops.
- The use of the large adaptive gains is important in the robustness of $\mathcal{L}_1$ adaptive control systems, as proven in Theorem 2.2.4 of [4].
- As we clarified earlier, while for some classes of systems $\mathcal{L}_1$ adaptive controller reduces to a simpler structure, this structure is ***not*** a filtered PI controller. Moreover, the authors have missed important references such as [15].

References

[1] P. Ioannou, K. S. Narendra, A. M. Annaswamy, S. Jafari, L. Rudd, R. Ortega, and J. Boskovic, "$\mathcal{L}_1$-adaptive control: Stability, robustness, and misperceptions," *IEEE Transactions on Automatic Control*, 2012, submitted.

[2] C. Cao and N. Hovakimyan, "Design and analysis of a novel $\mathcal{L}_1$ adaptive control architecture, Part I: Control signal and asymptotic stability," in *American Control Conference*, Minneapolis, MN, June 2006, pp. 3397–3402.

[3] ——, "Design and analysis of a novel $\mathcal{L}_1$ adaptive control architecture, Part II: Guaranteed transient performance," in *American Control Conference*, Minneapolis, MN, June 2006, pp. 3403–3408.

[4] N. Hovakimyan and C. Cao, *$\mathcal{L}_1$ Adaptive Control Theory*. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2010.

[5] B. D. O. Anderson, "Failures of adaptive control theory and their resolution," *Communications in Information and Systems*, vol. 5, no. 1, pp. 1–20, 2005.

[6] N. Hovakimyan, C. Cao, E. Kharisov, E. Xargay, and I. M. Gregory, "$\mathcal{L}_1$ adaptive control for safety-critical systems. guaranteed robustness with fast adaptation," *IEEE Control Systems Magazine*, October 2011.

[7] E. Kharisov, N. Hovakimyan, and K. J. Åström, "Comparison of several adaptive controllers according to their robustness metrics," in *AIAA Guidance, Navigation and Control Conference*, Toronto, Canada, August 2010, accepted for publication.

[8] A. S. Morse, "Global stability of parameter-adaptive control systems," *IEEE Transactions on Automatic Control*, vol. 25, no. 3, pp. 433–439, June 1980.

[9] G. Stein, "Respect the unstable," *IEEE Control Systems Magazine*, vol. 23, no. 4, pp. 12–25, August 2003.

[10] C. Cao and N. Hovakimyan, "Stability margins of $\mathcal{L}_1$ adaptive control architecture," *IEEE Transactions on Automatic Control*, vol. 55, no. 2, pp. 480–487, February 2010.

[11] ——, "Guaranteed transient performance with $\mathcal{L}_1$ adaptive controller for systems with unknown time-varying parameters: Part I," in *American Control Conference*, New York, NY, July 2007, pp. 3925–3930.

[12] J.-B. Pomet and L. Praly, "Adaptive nonlinear regulation: Estimation from the Lyapunov equation," *IEEE Transactions on Automatic Control*, vol. 37, no. 6, pp. 729–740, June 1992.

[13] R. Cristi, J. J. Kim, and J. H. Yang, "Robust $\mathcal{L}_1$ adaptive control," July 2010.

[14] R. V. Monopoli, "Model reference adaptive control with an augmentation error signal," *IEEE Transactions on Automatic Control*, vol. 19, no. 5, October 1974.

[15] E. Kharisov and N. Hovakimyan, "$\mathcal{L}_1$ adaptive output feedback controller for minimum phase systems," in *American Control Conference*, San Francisco, CA, June 2011.

[16] K. K. K. Kim, E. Kharisov, and N. Hovakimyan, "Filter design for $\mathcal{L}_1$ adaptive output–feedback controller," in *IEEE Conference on Decision and Control*, Orlando, FL, December 2011, accepted.

[17] D. Li, V. V. Patel, C. Cao, N. Hovakimyan, and K. A. Wise, "Optimization of the time-delay margin of $\mathcal{L}_1$ adaptive controller via the design of the underlying filter," in *AIAA Guidance, Navigation and Control Conference*, Hilton Head, SC, August 2007, AIAA-2007-6646.

[18] M. Naghnaeian, P. G. Voulgaris, and N. Hovakimyan, "On the robustness of $\mathcal{L}_1$ adaptive control with time varying perturbations & filter design," in *American Control Conference*, Montreal, Canada, June 2012, accepted.

[19] R. W. Beard, N. B. Knoebel, C. Cao, N. Hovakimyan, and J. S. Matthews, "An $\mathcal{L}_1$ adaptive pitch controller for miniature air vehicles," in *AIAA Guidance, Navigation and Control Conference*, Keystone, CO, August 2006, AIAA-2006-6777.

[20] J. Wang, V. Patel, C. Woolsey, N. Hovakimyan, and D. Schmale, "Adaptive control of a uav for aerobiological sampling," in *American Control Conference*, New York, NY, 2007, pp. 4660–4665.

[21] L. Techy, K. Reddy, C. Woolsey, C. Cao, and N. Hovakimyan, "Nonlinear control of a novel two-link pendulum," in *American Control Conference*, New York, NY, 2007, pp. 19–24.