Calhoun: The NPS Institutional Archive
DSpace Repository

Faculty and Researchers

Faculty and Researchers' Publications

2011

Experimental Validation of L1 Adaptive Control: The Rohrs Counterexample in Flight

Dobrokhodov, Vladimir; Kaminer, Isaac; Kitsios, Ioannis;
Xargay, Enric; Hovakimyan, Naira; Cao, Chengyu;
Gregory, Irene M.; Valavani, Lena

Journal of Guidance, Control, and Dynamics, Vol. 32, No. 5, September - October 2011
<https://hdl.handle.net/10945/45500>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun

Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

Experimental Validation of $\mathcal{L}_1$ Adaptive Control: The Rohrs Counterexample in Flight

Vladimir Dobrokhodov,^{*} Isaac Kaminer,[†] and Ioannis Kitsios[‡]

Naval Postgraduate School, Monterey, California 93943

Enric Xargay[§] and Naira Hovakimyan[¶]

University of Illinois at Urbana–Champaign, Urbana, Illinois 61801

Chengyu Cao^{**}

University of Connecticut, Storrs, Connecticut 06269

Irene M. Gregory^{††}

NASA Langley Research Center, Hampton, Virginia 23681

and

Lena Valavani^{‡‡}

Hellenic Space Systems, S.A., 154 51 Athens, Greece

DOI: 10.2514/1.50683

This paper presents flight-test results that examine the performance and robustness properties of an $\mathcal{L}_1$ control augmentation loop implemented onboard a small unmanned aerial vehicle. The framework used for in-flight control evaluation is based on the Rohrs counterexample, a benchmark problem presented in the early 1980s, to show the limitations of adaptive controllers developed at that time. Hardware-in-the-loop simulations and flight-test results confirm the ability of the $\mathcal{L}_1$ flight control system to maintain stability and predictable performance of the closed-loop adaptive system in the presence of general (artificially injected) unmodeled dynamics. The results demonstrate the advantages of $\mathcal{L}_1$ control as a robust adaptive control architecture with the potential of facilitating the transition of adaptive control into advanced flight control systems.

I. Introduction

ADAPTIVE flight control systems are seen as an appealing technology that may provide the opportunity to improve aircraft performance and reduce pilots' workload at challenging flight envelope conditions or in the event of severe failures and vehicle damage. However, several limitations have been identified that need to be addressed to enable the transition of adaptive control technologies into safety-critical aerospace applications. Incomplete overviews on fundamental deficiencies of adaptive control are provided in [1,2]. Additionally, discussions on open problems and certification challenges for adaptive flight control systems can be found in [3–5]. In particular, the key deficiencies of adaptive (flight) control systems

are 1) the lack of predictability in the closed-loop response, 2) the limited analysis framework for robustness and performance guarantees for closed-loop adaptive systems, and 3) the lack of systematic design guidelines to solve the tradeoff between adaptation, performance, and robustness. These limitations seem to be directly related to the asymptotic nature of the results obtained in the development of the theory of adaptive control over the years. In fact, when dealing with practical applications, features such as boundedness, ultimate boundedness, or even asymptotic convergence are weak properties for nonlinear adaptive feedback systems. Much stronger guarantees are needed. On one hand, performance requirements demand predictable and consistent response of the closed-loop system, dependent upon changes in system dynamics and reference signals. On the other hand, system uncertainty requires accurate quantification of the robustness and the stability margins of the feedback loop. Moreover, when it comes to the application of adaptive controllers in a real problem, the lack of analytical quantification of the relationship between the adaptation process, the transient response, and the robustness margins makes the design of such controllers an overly challenging problem, which is being commonly resolved by either computationally expensive Monte Carlo simulations or trial and error methods.

Several examples have been presented over the years to illustrate these limitations. The lack of robustness exhibited by model reference adaptive controllers was first identified in 1978 by Egardt [6] and later analyzed by Rohrs et al. [7,8]. In particular, Rohrs et al. constructed a counterexample where a first-order stable nominal plant with two (apparently harmless) highly damped unmodeled poles was experiencing instability when driven by a reference sinusoid at the phase crossover frequency or affected by unmeasurable output disturbances at any frequency. The lack of robustness of adaptive controllers has also been analyzed in robust control literature. In [9], Georgiou and Smith proved that a conventional parameter adaptive controller has zero robustness margin in the gap metric and can be destabilized by arbitrarily small perturbations in the gap. Moreover, the lack of transient characterization of asymptotically stable adaptive controllers was demonstrated by several

Presented as Paper 2009-6188 at the AIAA Guidance, Navigation, and Control Conference, Chicago, IL, 10–13 August 2009; received 9 May 2010; revision received 19 November 2010; accepted for publication 28 January 2011. Copyright © 2011 by V. Dobrokhodov, I. Kaminer, I. Kitsios, E. Xargay, N. Hovakimyan, C. Cao, I. M. Gregory, and L. Valavani. Published by the American Institute of Aeronautics and Astronautics, Inc., with permission. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0731-5090/11 and \$10.00 in correspondence with the CCC.

^{*}Research Assistant Professor, Department of Mechanical and Aerospace Engineering; vldobr@nps.edu. Senior Member AIAA.

[†]Professor, Department of Mechanical and Aerospace Engineering; kaminer@nps.edu. Senior Member AIAA.

[‡]Postdoctoral Research Fellow, Department of Mechanical and Aerospace Engineering, Cpt (HAF); ikitsios@nps.edu.

[§]Doctoral Student, Department of Aerospace Engineering; xargay@illinois.edu. Student Member AIAA.

[¶]Professor, Department of Mechanical Science and Engineering; nhovakim@illinois.edu. Associate Fellow AIAA.

^{**}Assistant Professor, Department of Mechanical Engineering; ccao@engr.uconn.edu. Member AIAA.

^{††}Senior Research Engineer, Dynamic Systems and Controls Branch; irene.m.gregory@nasa.gov. Associate Fellow AIAA.

^{‡‡}President, Hellenic Space Systems S.A.; valavani@mit.edu. Associate Fellow AIAA.

counterexamples that appeared in the literature [10–15]. These examples showed that the system output can have very poor transient behavior before ideal asymptotic convergence takes place.

In particular, the results and conclusions of Rohrs et al.'s counterexample presented in [7] (and in its journal version [8]) led to an ideological controversy and, as a consequence, robustness of adaptive controllers started to be investigated by other authors. A thorough explanation for the phenomena observed in Rohrs et al.'s simulations was provided in papers by Åström [16,17] and, more recently, by Anderson [1]. Other authors attempted to provide a solution to the problem of parameter drift with limited success, as the modifications they proposed were essentially compromising the adaptation process in varying degrees and, ultimately, could not always predictably prevent the parameter drift [18–23]. The basic idea of all the modifications was to limit the gain of the adaptation loop and to eliminate its integral action. Examples of these modifications are the σ modification [22] and the e modification [23]. A good survey on the topic can be found in [24]. While all of these modifications offered some improvement in terms of robustness, none of them provided the means to 1) quantify the transient response of the closed-loop adaptive system, 2) develop a framework for analysis of its performance and robustness characteristics, or 3) provide systematic guidelines for its design.

The $\mathcal{L}_1$ adaptive control theory [25], which appeared as a method for the design of robust adaptive control architectures using fast estimation schemes, precisely addresses these limitations by setting in place an architecture in which adaptation is decoupled from robustness. In these architectures, the speed of adaptation is limited only by the available hardware (computational power and high-frequency sensor noise), while the tradeoff between performance and robustness can be resolved via conventional methods from robust control. The architectures of $\mathcal{L}_1$ adaptive control theory have guaranteed transient performance and guaranteed robustness in the presence of fast adaptation, without introducing or enforcing persistence of excitation, without gain scheduling in the controller parameters, and without resorting to high-gain feedback.

In this paper, we present flight-test results that examine the performance and robustness properties of an $\mathcal{L}_1$ flight control augmentation system implemented onboard a small unmanned aerial vehicle (UAV). The framework used for in-flight validation is inspired by the Rohrs counterexample, which we extend here to a flight-test environment. The first-order nominal stable plant used in Rohrs et al.'s simulations [7,8] is replaced by a small UAV controlled by a commercial autopilot (AP). This nominal plant, consisting of the closed-loop UAV with its AP, is then augmented with an adaptive controller for improved angular-rate tracking capabilities in the presence of different artificially injected unmodeled dynamics. This flight-test setup allows for verification of the theoretical claims of $\mathcal{L}_1$ adaptive control and for demonstrating the advantages of $\mathcal{L}_1$ adaptive control as a robust adaptive control architecture with the potential of facilitating the transition of adaptive control into advanced flight control systems of general aviation.

The paper is organized as follows. Section II presents the flight-test setup adopted for in-flight control evaluation. Section III addresses the identification of the frequency response of the nominal plant consisting of the UAV with its AP and presents the two output-feedback adaptive algorithms tested in flight. Hardware-in-the-loop (HIL) experiments are used to tune these adaptive algorithms to achieve desired performance and robustness characteristics. Section IV extends the Rohrs counterexample to the flight-test environment and describes the results obtained in HIL simulations and flight tests. Finally, Sec. V summarizes the key results and contains the main conclusions.

II. Preliminaries of Study

A. Rapid Flight-Test Prototyping System

Recognizing the value of experimental evaluation of advanced flight control algorithms, the Naval Postgraduate School (NPS) team has developed the rapid flight-test prototyping system (RFTPS). The RFTPS consists of a testbed UAV equipped with a commercial AP

Fig. 1 Adaptive augmentation loop for turn-rate tracking tested by NPS.

(Piccolo Plus), an embedded computer running the research control algorithms in real time, and a ground control station for flight management and data monitoring and collection. In particular, the control algorithms are implemented on a MSM900BEV industrial PC104 computer using the xPC/RTW Target development environment and run at 100 Hz with the standard fixed-step ODE3 Bogacki–Shampine solver.^{§§} This embedded computer communicates with the AP over a full duplex serial link at 20 Hz. Also, the onboard avionics are augmented with a wireless communication link, which is used to (bidirectionally) exchange telemetry data in real time between the AP and the ground station, allowing for real-time control, tuning, and performance monitoring of the developed software. This system facilitates the real-time onboard integration of advanced control algorithms, and it provides the opportunity to design and conduct comprehensive flight-test programs to evaluate the robustness and performance characteristics of these algorithms. More details on the architecture of the developed flight-test system can be found in [26].

To demonstrate the benefits of $\mathcal{L}_1$ adaptive control, the commercial AP of the RFTPS has been augmented with the $\mathcal{L}_1$ output-feedback architectures presented in [27,28]. The $\mathcal{L}_1$ augmentation loop is introduced to enhance the turn-rate tracking capabilities of the commercial AP in the event of control surface failures and in the presence of significant environmental disturbances. The adaptive flight control architecture implemented on the RFTPS is represented in Fig. 1. Preliminary flight-test results of this setup were first reported in [29], where Dobrokhodov et al. analyzed stability and performance of the $\mathcal{L}_1$ adaptive control system in the presence of locked-in-place control surface failures. The results obtained demonstrated that the $\mathcal{L}_1$ augmented system provides fast recovery to sudden failures in one of the ailerons or in the rudder, while the unaugmented system would go unstable.

In the present paper, the RFTPS is modified to permit the artificial injection of unmodeled dynamics at the output of the closed-loop UAV and AP system (see Fig. 2). This setup, which is inspired by the Rohrs counterexample, provides the opportunity to evaluate the robustness and performance characteristics of flight control systems in the presence of different system uncertainties and disturbances. In particular, for the purpose of this study, three cases of unmodeled dynamics are considered. In the first case, a second-order transfer function introducing uncertainty both at high frequencies and inside the bandwidth of the UAV dynamics is injected. In the second case, we consider a very lightly damped second-order transfer function introducing uncertainty well beyond the bandwidth of the (rigid) UAV as an attempt to replicate the effects of flexible body modes. This second setup hence allows for identification of possible structural mode interactions between the UAV and the adaptive augmentation loop implemented onboard. Finally, the third scenario considers the injection of a locked-in-place control surface failure in the left aileron.

To verify that this setup replicates the phenomena observed in Rohrs et al.'s simulations [7,8] and that it appropriately reflects the undesirable and unpredictable interactions between the adaptation

^{§§}We notice that the base sampling time and the solver used for implementation of the control algorithms can be adjusted according to specific needs.

Fig. 2 Closed-loop adaptive system with artificially injected unmodeled dynamics and disturbances.

process and the closed-loop system dynamics, we first augment the onboard AP with a model reference adaptive control (MRAC) algorithm with similar robustness properties as the adaptive algorithm considered by Rohrs et al. In addition, we also implement some of the modifications developed to overcome the problem of parameter drift in conventional MRAC and demonstrate that these modifications in fact recover stability of the closed-loop adaptive system. We emphasize that the objective of these preliminary experiments is only to calibrate the flight-test setup and to verify correctness of the proposed framework. The flight-test setup is then used to evaluate the stability, performance, and robustness properties of the $\mathcal{L}_1$ augmentation loop and to verify the theoretical claims for $\mathcal{L}_1$ adaptive control architectures.

B. Bowditch–Lissajous Curves

To replicate Rohrs et al.'s simulations [7,8] with the flight-test setup introduced above, we first need to determine the frequency response of the nominal plant consisting of the UAV and its AP. To this end, we exploit the properties of Bowditch–Lissajous curves. A Bowditch–Lissajous curve is a curve described by the equations

$$x(t) = A \sin(\omega_x t - \phi_x) \quad y(t) = B \sin(\omega_y t - \phi_y)$$

which are parameterized by the amplitudes A and B , the frequencies ω_x and ω_y , and the phase shifts ϕ_x and ϕ_y . The shape of this two-dimensional curve is a function of its parameters and, in particular, it is closed if and only if the ratio ω_x/ω_y is rational. More details on Bowditch–Lissajous curves can be found in [30].

It is well known that, when a stable linear-time-invariant (LTI) system is driven by a sinusoidal input signal, its output response is also a sinusoidal signal at the same frequency but with different amplitude and phase shift, the last two being functions of the frequency of the input signal. Then, the curve generated by plotting the output of the LTI system against its input is an ellipse, which is a Lissajous curve for the special case of $\omega_x = \omega_y$. The eccentricity of the Lissajous curve is related to the phase shift between input and output, while the ratio of the projections of its body onto the y and x axes defines the magnitude of the frequency response.

These properties of the Lissajous curves are exploited in this paper for determination of the frequency response of the closed-loop UAV with its AP while being operated in its linear range. In addition, the Lissajous curves appear to be a useful tool to evaluate performance of adaptive flight control systems. In fact, the Lissajous curve generated by plotting the output response signal of the actual closed-loop adaptive system against the output response of the corresponding reference system provides a quantitative and accurate assessment of the performance of the adaptive control algorithm, eliminating any ambiguities in the noisy experimental data. Moreover, for the particular application considered in this paper, in which an inner-loop AP is augmented with an adaptive controller, the Lissajous curve obtained by plotting the adaptive control signal (which is the command signal sent to the AP; see Fig. 1) against the reference signal to be tracked offers very valuable information about the behavior of the adaptive controller. The quick and intuitive interpretation of Lissajous curves was particularly effective in monitoring

in real time the performance of the adaptive controllers during the flight-test experiments.

III. Frequency-Response Analysis and Adaptive Augmentation Loop Design

A. Frequency-Response Analysis of the Nominal Plant

As mentioned earlier, the extension of the Rohrs counterexample requires identification of the frequency response of the closed-loop UAV with its AP. This result will be useful later in the paper to determine the phase crossover frequency of the system with the artificially injected uncertainties, which is a necessary step to reproduce Rohrs et al.'s results [7,8]. At the same time, having a (linear) model of the nominal plant is also important for the design of the adaptive augmentation loops, as the architecture of adaptive output-feedback schemes usually depends on the structure of the plant to be controlled.

To determine the frequency response of the nominal plant, a sinusoidal signal was sent as a turn-rate command to the unaugmented AP. In particular, this reference signal was of the form

$$r_{cmd}(t) = A \sin(\omega t)$$

with tunable A amplitude and ω frequency. A set of Lissajous curves for the turn-rate response of the closed-loop UAV with the AP, together with their filtered least-mean-square estimates (in gray), is presented in Fig. 3. Table 1 summarizes the frequency-response results derived from the Lissajous curves estimates. These results correspond to HIL simulations.

An analysis of the frequency-response results of the nominal system and previous system identification work on the Rascal UAV (see [31]) shows that the turn-rate response of the nominal plant can be modeled in the linear range as the following second-order transfer function with two underdamped dominant poles and relative degree $n^* = 1$:

$$G_p(s) = \frac{(1/0.5)s + 1}{(1/0.55^2)s^2 + (2 \cdot 0.8/0.55)s + 1} \quad (1)$$

In fact, a first-order system adequately captures the turn-rate dynamics of the bare airframe, while the AP, which has a proportional-integral structure, ensures a unity dc gain of the turn-rate closed-loop system and introduces a minimum-phase zero.

B. Adaptive Augmentation Algorithms

In this paper, the AP mounted onboard the UAV is augmented with an adaptive output-feedback controller that modifies the turn-rate reference signal based on actual turn-rate measurements and sends the augmented command to the AP. As explained earlier, the main objective of wrapping an adaptive augmentation around the AP is to

Fig. 3 Lissajous curves $r(r_{cmd})$ of (unaugmented) nominal plant for various frequencies.

Table 1 Frequency response of closed-loop UAV with its AP

Frequency, rad/s	0.1	0.5	0.75	1.0	1.5	2.0
Gain, dB	0.1	-0.4	-0.63	-2.5	-2.9	-6.8
Phase, deg	-5.4	-33.1	-45.85	-54.4	-78.7	-78.9

analyze whether turn-rate tracking performance and aircraft safety can be improved in the event of control surface failures and vehicle damage.

In this section, we present different adaptive algorithms that have been implemented onboard the RFTPS and tested both in HIL simulations and in-flight experiments. First, we introduce a conventional output-feedback MRAC algorithm with properties similar to the adaptive controller used in the Rohrs counterexample. In addition, we introduce some of the adaptive law modifications developed over the years to overcome the problem of parameter drift in conventional MRAC. These adaptive algorithms, with and without modifications, are used to check the validity of the flight-test setup and to verify that the phenomena observed in Rohrs et al.'s simulations [7,8] can be replicated. Finally, we also present the $\mathcal{L}_1$ output-feedback control architecture implemented on the RFTPS, which has been proven to enhance the angular-rate tracking capabilities of commercial APs [29,32]. For the design of both the MRAC and the $\mathcal{L}_1$ augmentation algorithms, we will assume that the closed-loop UAV and its AP is represented by a single-input/single-output uncertain second-order transfer function with relative degree 1 and known sign of the high-frequency gain, and we will consider the system in Eq. (1) as the nominal model of the plant. Then, the augmentation loops will be tuned so that the nominal augmented closed-loop system achieves a similar level of tracking performance for sinusoidal reference signals in the low-frequency range, with a similar time-delay margin at the input of the plant.

1. Model Reference Adaptive Control Augmentation Algorithm

Several MRAC algorithms are available in the literature with similar robustness properties as the controller used in the Rohrs counterexample. References [33,34] provide a good overview of this class of adaptive controllers. For the purpose of this paper, we use the algorithm presented in [34] (Sec. 6.4). The structure of this controller and some of the assumptions that need to be verified are introduced next. The reader is referred to [34] for more details on this algorithm. We also provide details on the implementation of the algorithm for the particular application at hand and show HIL simulation results to demonstrate its performance capabilities.

a. Direct Model Reference Adaptive Control with Unnormalized Adaptive Laws. This algorithm assumes a single-input/single-output plant of the following form:

$$G_p(s) = k_p \frac{Z_p(s)}{R_p(s)}$$

where $Z_p(s)$ and $R_p(s)$ are unknown monic polynomials; thus, k_p represents the unknown high-frequency gain. The plant is assumed to be minimum phase, while an upper bound n on the number of poles, the relative degree of the plant n^* , and the sign of the high-frequency gain are assumed to be known.

The reference model that describes the desired dynamics is given by

$$G_m(s) = k_m \frac{Z_m(s)}{R_m(s)}$$

where $Z_m(s)$ and $R_m(s)$ are monic Hurwitz polynomials, while k_m is a constant. The reference model $G_m(s)$ is assumed to be strictly positive real (SPR) and with the same relative degree n^* as the plant $G_p(s)$.

The control law $u(t)$ that solves the model reference problem can be formulated using the following state-space realization:

Fig. 4 MRAC augmentation loop.

$$\begin{aligned}\dot{\omega}_1(t) &= F\omega_1(t) + gu(t), & \omega_1(0) &= 0 \\ \dot{\omega}_2(t) &= F\omega_2(t) + gy(t), & \omega_2(0) &= 0 \\ u(t) &= \theta^T(t)\omega(t)\end{aligned}\quad (2)$$

where the Hurwitz matrix $F \in \mathbb{R}^{(n-1) \times (n-1)}$ and the vector $g \in \mathbb{R}^{n-1}$ are degrees of freedom available to the controller designer, $\omega_1(t)$, $\omega_2(t) \in \mathbb{R}^{n-1}$ are internal regression signals, and $\theta(t) \in \mathbb{R}^{2n}$ is the vector of parameter estimates. Let $\omega(t) \in \mathbb{R}^{2n}$ be given by

$$\omega(t) = [\omega_1^T(t), \omega_2^T(t), y(t), r(t)]^T$$

where $y(t)$ and $r(t)$ represent the output of the plant and the reference signal, respectively. Then, the adaptation law for $\theta(t)$ is given by

$$\dot{\theta}(t) = -\Gamma e(t)\omega(t)\text{sgn}\left(\frac{k_m}{k_p}\right), \quad \theta(0) = \theta_0 \quad (3)$$

where $\Gamma > 0$ and $e(t) = y(t) - y_m(t)$, with $y_m(t)$ being the output of the reference model $G_m(s)$ to the reference signal $r(t)$. The block diagram for this MRAC algorithm is presented in Fig. 4. For details on the choice of the controller parameters F , g , and Γ , the reader is again referred to [34].

This formulation of the MRAC algorithm guarantees the following:

1) All signals of the closed-loop adaptive system are bounded, and the tracking error $e(t)$ converges to zero asymptotically for any bounded reference signal $r(t)$.

2) If the reference signal $r(t)$ is sufficiently rich, $\dot{r}(t)$ is bounded, and the polynomials $Z_p(s)$ and $R_p(s)$ are relatively coprime, then the parameter error and the tracking error converge exponentially to zero.

To improve the robustness properties of this algorithm, we modify the MRAC adaptation law with the 1) projection operator,

$$\dot{\theta}(t) = \Gamma \text{Proj}(\theta(t), -e(t)\omega(t)\text{sgn}(k_m/k_p)), \quad \theta(0) = \theta_0$$

where Proj represents the projection operator defined over a given compact set Θ .[¶] 2) σ modification,

$$\dot{\theta}(t) = -\Gamma e(t)\omega(t)\text{sgn}(k_m/k_p) - \sigma\theta(t), \quad \theta(0) = \theta_0$$

where σ is a tunable positive parameter; and 3) e modification,

$$\dot{\theta}(t) = -\Gamma e(t)\omega(t)\text{sgn}(k_m/k_p) - \gamma\theta(t)|e(t)|, \quad \theta(0) = \theta_0$$

where γ is a tunable positive parameter.

b. Model Reference Adaptive Control Augmentation Loop Design. As shown in Sec. III.A, the closed-loop UAV with the AP can be modeled in the low-frequency range as a second-order model

[¶]See the Appendix for the definition of the projection operator used in the adaptive laws.

of relative degree $n^* = 1$. Then, according to the design constraints of the MRAC algorithm presented above, we can choose the reference model to be a stable first-order system of the form

$$G_m(s) = \frac{m}{s + m}, \quad m > 0$$

which, at the same time, is an SPR transfer function. In fact, one can check that this reference model verifies all the assumptions stated in the description of the MRAC algorithm introduced in the previous section. The bandwidth of this reference model was chosen to be 1.5 rad/s, which is similar to the -3 dB bandwidth of the nominal airplane with the AP, as can be seen in Table 1.

For the design of the MRAC algorithm, we will assume that the upper bound on the number of poles in the plant is $n = 2$. With this assumption, both $\omega_1(t)$ and $\omega_2(t)$ are just scalar signals, while the parameters F and g are scalar positive constants. Therefore, the vector $\omega(t)$ belongs to $\mathbb{R}^4$ and we will have a total of four adaptive parameters:

$$\theta(t) = [\theta_1(t), \theta_2(t), k_y(t), k_r(t)]^\top$$

where $\theta_1(t)$, $\theta_2(t)$, $k_y(t)$, and $k_r(t)$ are scalar parameters. The matrix of adaptive gains will be chosen to be a diagonal 4×4 matrix with positive entries.

The tunable parameters of the MRAC algorithm have been adjusted by trial and error in HIL experiments with the objective of

achieving a satisfactory convergence rate for the parameters without compromising the robustness of the closed-loop adaptive system. The values of the parameters are given as follows:

$$F = -2, \quad g = 1, \quad \Gamma = \text{diag}([5, 10, 2, 2.5])$$

while the initial parameter estimate was set to

$$\theta_0 = [0, 0, 0, 1]$$

For this set of parameters, the time-delay margin (defined at the input of the system) obtained in simulation at the speed of 22 m/s and altitude of 550 m is $\tau^* \approx 0.15$ s.

Figure 5 shows the results of one of the experiments in the HIL environment. In particular, it shows the response of the closed-loop adaptive system with the MRAC augmentation loop to a set of biased sinusoidal reference signals at different frequencies. The reference signals used for the flight-test experiments are of the form

$$r_{\text{cmd}}(t) = b + A \sin(\omega t)$$

with tunable b bias, A amplitude, and ω frequency. Introducing a bias term b results in the airplane orbiting continuously in a bounded airspace box, which makes the flight experiment significantly easier to perform; the value of this bias term is chosen small enough to allow a reasonable flexibility in choosing the amplitude A without internally saturating the AP. For consistency, we also use biased

Fig. 5 MRAC closed-loop response to a biased sinusoidal reference signal at a–b) $\omega = 0.1$ rad/s, c–d) $\omega = 0.3$ rad/s, and e–f) $\omega = 0.5$ rad/s.

sinusoidal reference signals in HIL simulations. As one can see, the MRAC algorithm is able to asymptotically track the output of the reference system, denoted in this paper by r_m , responding to the different sinusoidal reference signals. Figures 5a, 5c, and 5e present the Lissajous curves generated by plotting the actual turn rate of the UAV against the output of the reference system used in the implementation of the adaptive controller. Note that these curves are close to straight lines of slope 1, which demonstrates that the tracking objective of MRAC is (asymptotically) achieved.

Together with the tunable parameters of the MRAC algorithm, the different modifications of the adaptive laws were also adjusted in HIL. The objective was to improve the robustness of the standard MRAC algorithm without sacrificing significantly the performance of the closed-loop adaptive system for reference signals in the low-frequency range. The parameters chosen for the different modifications are detailed next:

1) For the projection operator, the bounds for the adaptive parameters were chosen as follows:

$$\theta_1(t), \theta_2(t), k_y(t) \in [-0.3, 0.3], \quad \text{and} \quad k_r(t) \in [0.1, 2]$$

2) For the σ modification, the value of the tunable parameter σ was chosen to be $\sigma = 0.075$.

3) For the e modification, the value of the tunable parameter γ was chosen to be $\gamma = 0.25$.

2. $\mathcal{L}_1$ Augmentation Algorithm

This section provides an overview of the $\mathcal{L}_1$ adaptive output-feedback controller for systems of unknown dimension in the presence of unmodeled dynamics and time-varying uncertainties. The reader is referred to [27] for a more detailed explanation of this architecture, as well as for the main results and their proofs. The section also includes HIL simulation results to illustrate the performance characteristics of this algorithm.

a. $\mathcal{L}_1$ Adaptive Output-Feedback Controller for First-Order Reference Systems. We start by considering the system

$$y(s) = G_p(s)[u(s) + z(s)], \quad y(0) = 0 \quad (4)$$

where $u(t) \in \mathbb{R}$ is the system's input, $y(t) \in \mathbb{R}$ is the system's output, $G_p(s)$ is assumed to be an unknown strictly proper transfer function, $z(s)$ is the Laplace transform of the time-varying uncertainties and disturbances $d(t) = f[t, y(t)]$, and $f: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is an unknown map, subject to the following assumptions.

Assumption 1: There exist constants $L > 0$ and $L_0 > 0$ such that the following inequalities

$$|f(t, y_1) - f(t, y_2)| \leq L|y_1 - y_2|, \quad |f(t, y)| \leq L|y| + L_0$$

hold uniformly in $t \geq 0$. We note that the numbers L and L_0 can be arbitrarily large.

The control objective is to design an output-feedback controller $u(t)$ such that the system output $y(t)$ tracks the given bounded piecewise continuous reference signal $r(t)$ following a desired reference model $M(s)$. In this section, we consider a first-order reference system; that is,

$$M(s) = \frac{m}{s + m}, \quad m > 0 \quad (5)$$

which is the same reference model that was implemented for the MRAC algorithm.

To provide an intuitive explanation behind the main idea of the $\mathcal{L}_1$ augmentation loop used in this paper, we note that the system in Eq. (4) can be rewritten in terms of the desired system $M(s)$ as

$$y(s) = M(s)[u(s) + \sigma(s)], \quad y(0) = 0 \quad (6)$$

where the uncertainties due to $G_p(s)$ and $z(s)$ are lumped into the signal $\sigma(s)$, which is defined as

$$\sigma(s) = \frac{[G_p(s) - M(s)]u(s) + G_p(s)z(s)}{M(s)} \quad (7)$$

The philosophy of the $\mathcal{L}_1$ output-feedback controller is to obtain an estimate of the uncertain signal $\sigma(t)$ and define a control signal that compensates for the effects of this uncertainty within the bandwidth of a low-pass filter $C(s)$ introduced in the control channel. This filter guarantees that the output of the $\mathcal{L}_1$ controller stays in the low-frequency range in the presence of fast adaptation and large reference inputs, leads to separation between adaptation and robustness, and defines the tradeoff between performance and robustness. Adaptation is based on the projection operator, ensuring boundedness of the adaptive parameters by definition [35], and uses the output of an output predictor to update the estimate of the uncertainty $\hat{\sigma}(t)$. This output predictor is defined to have the same structure as the system in Eq. (6) but using the estimate $\hat{\sigma}(t)$ instead of $\sigma(t)$ itself, which is unknown. The $\mathcal{L}_1$ adaptive control architecture is represented in Fig. 6, and its elements are introduced below.

For the output predictor, we consider the following output predictor:

$$\dot{\hat{y}}(t) = -m\hat{y}(t) + m[u(t) + \hat{\sigma}(t)], \quad \hat{y}(0) = 0 \quad (8)$$

where $\hat{\sigma}(t)$ is the adaptive estimate.

For the adaptation law, the adaptation of $\hat{\sigma}(t)$ is defined as

$$\dot{\hat{\sigma}}(t) = \Gamma_c \text{Proj}[\hat{\sigma}(t), -\tilde{y}(t)], \quad \hat{\sigma}(0) = 0 \quad (9)$$

where $\tilde{y}(t) = \hat{y}(t) - y(t)$ is the output-prediction error signal, $\Gamma_c > 0$ is the adaptation rate subject to a computable lower bound, and Proj denotes the projection operator.

For the control law, the control signal is generated as

$$u(s) = r(s) - C(s)\hat{\sigma}(s) \quad (10)$$

where $C(s)$ is a strictly proper low-pass filter with $C(0) = 1$.

The complete $\mathcal{L}_1$ output-feedback controller consists of Eqs. (8–10) subject to the following stability conditions: the design of $C(s)$ and $M(s)$ must ensure that

$$H(s) = \frac{G_p(s)M(s)}{C(s)G_p(s) + [1 - C(s)]M(s)} \quad (11)$$

is stable and the following $\mathcal{L}_1$ -norm sufficient condition holds:

$$\|H(s)(1 - C(s))\|_{\mathcal{L}_1} L < 1 \quad (12)$$

Then, it has been proven that the control signal $u(t)$ and the output $y(t)$ of the closed-loop adaptive system with the $\mathcal{L}_1$ adaptive controller in Eqs. (8–10) track both in transient and steady state the input and output of an auxiliary closed-loop reference system, which is defined as the ideal nonadaptive version of the $\mathcal{L}_1$ adaptive

Fig. 6 $\mathcal{L}_1$ adaptive augmentation loop.

controller assuming perfect knowledge of uncertainties. We refer to [27] for a detailed explanation of these results and a complete derivation of their proofs.

b. $\mathcal{L}_1$ Augmentation Loop Design. The $\mathcal{L}_1$ augmentation scheme is tuned to achieve a similar level of performance as the MRAC algorithm for reference signals in the low-frequency range and, at the same time, have similar robustness characteristics. In particular, the reference model was chosen to be exactly the same as for the MRAC algorithm, i.e., a stable first-order system with bandwidth 1.5 rad/s. The adaptive gain and the projection bounds were set to 30,000 and ± 1 rad/s, while the low-pass filter was chosen to be a first-order system with unity dc gain and 0.6 rad/s of bandwidth. This design guarantees that, for the nominal plant in Eq. (1), the transfer function $H(s)$ in Eq. (11) is stable, and the $\mathcal{L}_1$ -norm sufficient condition in Eq. (12) holds for uncertainties with a Lipschitz constant $L \leq 0.9$. The adaptation gain was set high to provide satisfactory performance of the estimation algorithm (notice that the performance bounds in $\mathcal{L}_1$ adaptive control theory are inverse proportional to the square root of the adaptation rate; see [25]). Then, considering an ODE3 solver for the real-time implementation of the adaptive control algorithms, the base sampling time of the PC104 computer was set to 10 ms in order to accommodate the fast estimation loop and ensure numerical stability of the real-time code. This choice for the base sampling time also avoids the computational overload of the control algorithm, as the average task execution time of the entire control code (including the adaptive augmentation loop,

the software interfacing loop, and the overhead required by the real-time operating system) is less than 1.2 ms. The choice of the low-pass filter and the projection bounds provides an adequate propagation of measurement noise into the control signal. Moreover, the design of the low-pass filter ensures a satisfactory level of performance with a time-delay margin (in simulation) of $\tau^* \approx 0.15$ s at the same flight condition of 22 m/s and 550 m.

Figure 7 shows the results from one of the HIL experiments conducted to tune the $\mathcal{L}_1$ adaptive augmentation loop. In particular, it shows the response of the closed-loop adaptive system with the $\mathcal{L}_1$ augmentation loop to a set of biased sinusoidal reference signals at different frequencies. The $\mathcal{L}_1$ controller is able to perfectly track the output of the reference system to sinusoidal reference signals in the low-frequency range (≤ 0.1 rad/s), as the Lissajous curve in Fig. 7a illustrates. Moreover, one can observe in Fig. 7b that the contribution of the $\mathcal{L}_1$ controller $-C(s)\hat{\sigma}(s)$ to the reference signal r_{cmd} is almost zero during this experiment, and it only becomes significant at high turn rates, which are characterized by high bank-angle attitudes with nonlinear dynamics. This indicates that the AP mounted onboard the UAV is able to precisely track reference signals in the low-frequency range and in the linear range of the UAV, while the adaptive controller helps the AP only when needed. Furthermore, Figs. 7c–7f show that, as the frequency of the reference signal increases beyond the bandwidth of the low-pass filter $C(s)$, the eccentricity of the Lissajous curves decreases, which reveals a degradation in tracking performance, consistent with the theory of $\mathcal{L}_1$ adaptive control.

Fig. 7 $\mathcal{L}_1$ closed-loop response to biased sinusoidal reference signal at a–b) $\omega = 0.1$ rad/s, c–d) $\omega = 0.5$ rad/s, and e–f) $\omega = 1.0$ rad/s.

IV. Rohrs Counterexample in Flight

In this section, we describe the results obtained both in HIL simulations and flight tests for the different experiments proposed in Sec. II. In particular, we first verify that the phenomena observed in Rohrs et al.'s simulations [7,8], illustrating the limitations of conventional MRAC algorithms can be reproduced by the flight-test setup proposed in this paper. We also verify that the adaptive law modifications introduced to solve the problem of parameter drift in conventional MRAC algorithms in fact improve the robustness properties of the MRAC algorithm implemented onboard the small UAV. Specifically, results are only shown for the e modification, but similar results have been obtained for the MRAC algorithm with projection operator and with σ modification. We emphasize again that the objective of these first preliminary steps in the experiments is only to verify correctness of the proposed framework. Finally, we analyze the stability and performance characteristics of the closed-loop system with the $\mathcal{L}_1$ augmentation loop.

A. Hardware-in-the-Loop Simulations

1. Second-Order Unmodeled Dynamics

In this section, we artificially introduce at the output of the nominal plant (closed-loop UAV and AP) a second-order system representing unmodeled dynamics, similar to the Rohrs counterexample. In his counterexample, Rohrs et al. considered the case of very well damped unmodeled dynamics at high frequencies to show that, even in the presence of apparently harmless uncertainties, the stability of adaptive controllers was not guaranteed [7,8]. In this paper, nevertheless, we consider the case of more challenging uncertainties. In particular, we choose a low damped second-order system with natural frequency approximately equal to the bandwidth of the plant; that is,

$$\Delta(s) = \frac{\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2}$$

Fig. 8 Lissajous curve for nominal system with unmodeled dynamics driven by biased sinusoidal reference signal at phase crossover frequency ($\omega = 1.6$ rad/s).

Fig. 9 MRAC: closed-loop response in presence of second-order unmodeled dynamics to biased sinusoidal reference signals at different frequencies.

with $\omega_n = 1.5$ rad/s and $\zeta = 0.45$. With the addition of this second-order system, the new plant has the phase crossover frequency approximately at 1.6 rad/s. Figure 8 shows the Lissajous curve of the system output plotted against the reference signal for the cascaded system in response to the biased sinusoidal reference signal at 1.6 rad/s. As one can observe, the Lissajous curve exhibits a phase shift of -180° , which confirms that 1.6 rad/s corresponds to the phase crossover frequency of the cascaded system.

At this point, having identified the phase crossover of the system, we can extend Rohrs et al.'s results [7,8] to the flight-test setup considered in this paper. To this end, we consider the MRAC algorithm described in Sec. III. Figure 9 shows the response of the plant with the artificially injected second-order unmodeled dynamics to biased sinusoidal reference signals at different frequencies. The closed-loop adaptive system is able to asymptotically track the output of the reference model for reference signals in the low-frequency range ($\omega = 0.1$ rad/s and $\omega = 0.5$ rad/s). However, when the system is driven by a reference signal at $\omega = 1$ rad/s, bursting takes place ($t \in [330 \text{ s}, 350 \text{ s}]$). During this interval of time, the MRAC algorithm generates adaptive command signals with amplitudes reaching ± 100 deg/s, and the turn-rate tracking performance becomes very poor. Note that, for safety reasons, the AP has internal saturations that limit the commands sent to the UAV, which in this case prevents the aircraft from entering an abnormal flight condition. This temporary instability can be clearly observed in Figs. 9a and 9b. Also, Fig. 9c shows that there is an abrupt change in the adaptive parameter $\theta_1(t)$, which converges to a new value that stabilizes the system and achieves asymptotic tracking of the output of the reference model. Finally, when the system is driven by a reference signal at approximately the phase crossover frequency, $\omega = 1.6$ rad/s, parameter drift takes place and the closed-loop system with the MRAC algorithm becomes unstable.

These results are not new, and they just replicate in a flight-test environment the well known limitations of conventional MRAC algorithms shown by Rohrs et al. [7,8]. As we previously mentioned, several modifications of the adaptive laws were developed during the last 30 years to improve the robustness of MRAC architectures. Figure 10 shows the results of the same experiment obtained for the MRAC algorithm with the e modification. It can be seen that the parameters stay bounded (Fig. 10c) and the system remains stable for the whole experiment. However, as expected, significant degradation in tracking performance is observed. This degradation is evident in Fig. 10a, which presents Lissajous curves far from the desired straight line with slope 1, thus indicating that the output of the plant barely tracks the output of the reference model even for reference signals in the low-frequency range. It is important to note that further tuning of the MRAC algorithm and its modifications can be done, which might result in improved tracking performance. However, fine tuning will not change the nature of MRAC algorithms, for which there are no transient performance and robustness guarantees.

These results confirm that the proposed flight-test setup is able to replicate the adverse interactions between the adaptation process and the closed-loop system dynamics, verifying correctness of the framework developed in this paper. Next, we present the results

Fig. 10 MRAC with e modification: closed-loop response in presence of second-order unmodeled dynamics to biased sinusoidal reference signals at different frequencies.

obtained for the same experiment with the $\mathcal{L}_1$ augmentation loop (Fig. 11). As can be seen, the closed-loop adaptive system is able to precisely track the output of the reference system to the reference signal at $\omega = 0.1$ rad/s. As the frequency of the reference signal increases and goes beyond the bandwidth of the low-pass filter $C(s)$, the tracking performance degrades as expected. Figure 11a illustrates the consistency in the response of the closed-loop adaptive system, which traces very neat trajectories in the Lissajous plot. Moreover, Fig. 11c shows that the contribution of the $\mathcal{L}_1$ controller does not exhibit high-frequency content and remains within reasonable bounds, well inside the projection bounds for the (only) adaptive parameter $\hat{\sigma}(t)$. We emphasize that fast adaptation is the key for consistent input and output system response with desired performance in the low-frequency range, while the low-pass filter guarantees a desired level of robustness in the presence of this fast adaptation. The ability to adapt fast without sacrificing robustness of the closed-loop adaptive system is the main difference between $\mathcal{L}_1$ adaptive control and MRAC with or without its adaptive law modifications.

2. Structural Mode Interactions

In this section, we extend the philosophy of the previous experiment to the case of high-frequency unmodeled dynamics with very low damping. This experiment has, as a goal, to reproduce the effects of bending modes in an airplane and investigate the possible interaction between adaptive control and aircraft flexible dynamics. To this end, we choose a very low damped second-order transfer function with natural frequency equal to approximately seven times the bandwidth of the plant; that is,

$$\Delta(s) = \frac{\omega_n^2}{s^2 + 2\zeta\omega_n s + \omega_n^2}$$

with $\omega_n = 10$ rad/s and $\zeta = 0.01$ and, at the same time, we introduce a disturbance $d(t)$ in the feedback signal at exactly $\omega = 10$ rad/s:

$$d(t) = 20 \sin(10t) \text{ deg/s}$$

Fig. 11 $\mathcal{L}_1$: closed-loop response in presence of second-order unmodeled dynamics to biased sinusoidal reference signals.

Fig. 12 MRAC. Closed-loop response in the presence of bending modes to a biased sinusoidal reference signal at $\omega = 0.5$ rad/s.

Figure 12 shows the response of the system with the conventional MRAC algorithm to a biased sinusoidal reference signal at $\omega = 0.5$ rad/s. Initially, we wait for the MRAC parameters to converge, so that the output of the system tracks the output of the reference model. At $t = 50$ s, the dynamics representing the bending mode are injected, and the system runs in this configuration for approximately 30 s. The injection of the bending mode introduces very small oscillations at the output of the plant, probably due to the presence of noise. We can see that initially the adaptive parameter $k_y(t)$ is close to zero, which implies that these oscillations at the output of the plant are well attenuated inside the adaptive controller and do not propagate to the control signal that the MRAC augmentation loop sends to the AP. Therefore, the bending mode is not excited and, during these 30 s, the MRAC algorithm is able to track the output of the reference model without exciting the bending mode. At approximately $t = 120$ s, we inject the disturbance in the feedback signal. One can see that there is a 10 s transient phase with high-frequency oscillations at the output of the plant. These initial oscillations damp down and MRAC recovers a similar level of performance as the one before the injection of the disturbance. A similar result is obtained when the disturbance is disengaged at $t = 180$ s. However, when we inject the disturbance for a second time at $t = 200$ s, the adaptive parameters start drifting immediately. In particular, we see that the parameter $k_y(t)$ grows significantly fast, which implies that the high-frequency content at the output of the plant is amplified by the MRAC algorithm and sent to the AP, resulting in an unstable closed-loop adaptive system. This experiment illustrates that conventional MRAC is not able to cope with lightly damped high-frequency unmodeled dynamics, as it generates a counteracting control signal that tries to cancel the disturbances by fighting the attenuating capabilities of the AP. The result is an unstable closed-loop system that increasingly excites the bending mode.

When the MRAC algorithm is implemented with the e modification, the system retains stability along the whole experiment (Fig. 13). Nevertheless, the slow adaptation in the parameters due to a deficient adaptive law results in poor high-frequency disturbance

attenuation characteristics. The time segments $t = [80, 150]$ s and $t = [150, 220]$ s explicitly illustrate this phenomenon: as the bending mode and the disturbance are injected, the adaptive parameter $k_y(t)$ abruptly changes in an attempt to cancel the disturbance artificially injected in the feedback signal. As a result, this high-frequency disturbance is not sufficiently attenuated inside the adaptive augmentation loop, which leads to significant oscillations at the output of the plant.

Next, Fig. 14 presents the results obtained with the $\mathcal{L}_1$ controller. At $t = 50$ s the bending mode is injected, which (similar to MRAC) results in very small amplitude oscillations at the output of the system due to the presence of noise. During this initial phase, the contribution of the $\mathcal{L}_1$ controller stays in the low-frequency range. The disturbance is injected at $t = 108$ s. One can see that the $\mathcal{L}_1$ controller generates small oscillations in the control channel, which causes a 5 s transient at the output of the system with oscillations of approximately 1.5 Hz. These oscillations in the contribution of the $\mathcal{L}_1$ controller are due to the fact that the projection bounds in the adaptive law are too tight to handle the initial transient produced by the disturbance. Figure 15 shows that the uncertainty estimate $\hat{\sigma}(t)$ hits the bounds of projection, and as a consequence, the error between the output of the predictor $\hat{r}(t)$ is not able to track the output of the actual turn rate $r(t)$. These oscillations could be easily avoided by increasing the bounds of the projection operator. In the experiment, after these 5 s of transient, the oscillations damp down and the $\mathcal{L}_1$ adaptive controller recovers its initial performance. The presence of the low-pass filter is critical in attenuating the high-frequency content in the feedback signal and guaranteeing that the $\mathcal{L}_1$ contribution is within the low-frequency range. The attenuating capabilities of the $\mathcal{L}_1$ controller do not lead to an adverse interaction with the AP; thus, the attenuating properties of the nominal inner loop are preserved and the bending mode is not excited.

3. Control Surface Failures

In this section, instead of artificially introducing unmodeled dynamics at the output of the plant, we will consider the case of a

Fig. 13 MRAC with e modification: closed-loop response in presence of bending modes to biased sinusoidal reference signal at $\omega = 0.5$ rad/s.

Fig. 14 $\mathcal{L}_1$: closed-loop response in presence of bending modes to biased sinusoidal reference signal at $\omega = 0.5$ rad/s.

Fig. 15 $\mathcal{L}_1$ saturation in adaptive law.

Fig. 16 AP: response of nominal UAV with its AP to biased sinusoidal reference signal in event of left aileron failure.

sudden control surface failure affecting the nominal plant. In particular, the left aileron will get locked at -6 deg, and then, after some time, the aircraft will recover full lateral control.

First we show that the nominal plant with the AP is able to recover from the failure and track the reference signal with the remaining control authority.^{***} Figure 16 shows the response of the plant to a biased sinusoidal reference signal at $\omega = 0.5$ rad/s. The failure is injected at $t = 375$ s, and the AP takes around 20 s to recover from the failure, which is the time that the integrators of the AP need to readjust their contribution. The Lissajous curve shows that, in the presence of the failure, the AP is able to recover almost full performance with respect to the unimpaired healthy UAV.

Next, we consider the case of the conventional MRAC algorithm. Figure 17 shows the response of the closed-loop adaptive system. In the figure, one can see that initially the MRAC algorithm is able to track perfectly the output of the reference system. When the failure is introduced at $t = 190$ s, the MRAC algorithm is able to recover full performance in around 20 s and keep tracking the output of the reference system perfectly using the remaining control authority. The transient response is smooth and, as the integrators of the AP readjust, the command from the adaptive augmentation loop to the AP converges to the nominal adaptive contribution for the healthy UAV. However, when the UAV recovers full control authority ($t = 290$ s), the adaptive parameters drift and the closed-loop system with the MRAC algorithm becomes unstable. For the sake of clarity, we did not add to the Lissajous curve the third phase in which the UAV recovers full lateral control authority and the system becomes unstable.

Again, this result is not new. The only objective of this experiment is to verify that the same instability results could be obtained in the event of a control surface failure. Figure 18 shows that parameter drift can be avoided with the introduction of the e modification at the price of degraded performance. In this experiment, the aileron failure is introduced at $t = 125$ s, and the initial performance is recovered in about 35 s. Then, when the UAV regains full control authority at

$t = 250$ s, the UAV remains stable and eventually recovers the original performance after a transient of 50 s.

Figure 19 presents the results obtained for this experiment when the $\mathcal{L}_1$ controller is wrapped around the AP. The failure is introduced at $t = 100$ s and the $\mathcal{L}_1$ adaptive controller takes 15 s to regain the original level of performance using the remaining control authority. Note that even though the UAV has recovered the desired performance in 15 s, the integrators inside the AP keep readjusting during approximately 15 more seconds; therefore, the contribution of the $\mathcal{L}_1$ adaptive controller also has to readjust accordingly. The Lissajous curve shows that the $\mathcal{L}_1$ adaptive controller is able to maintain the initial phase shift of the output of the plant with respect to the output of the reference system, while the gain slightly increases for the impaired UAV. When the UAV regains full control authority at $t = 170$ s, the closed-loop adaptive system is able to stabilize the plant and achieve desired performance in 15 s. After 15 more seconds, when the integrators inside the AP readjust, the adaptive command of the $\mathcal{L}_1$ controller to the AP is the same as it was for the initial healthy UAV. Note that the transient characteristics of the closed-loop adaptive system when 1) the failure is introduced and 2) the UAV regains full control authority are similar, resulting in a predictable response of the closed-loop system.

To conclude this experiment, Fig. 20 shows a set of three Lissajous figures obtained by plotting the adaptive command sent to the AP against the reference signal. These figures illustrate the way the three different adaptive controllers modify the biased sinusoidal reference signal to achieve the desired tracking performance. In Fig. 20a, one can see that some oscillations appear in the command that the MRAC algorithm sends to the AP for the impaired UAV and, eventually, the system becomes unstable when full control authority is recovered. When the MRAC algorithm is implemented with the e modification, the contribution of the adaptive controller does not change significantly during the whole experiment (see Fig. 20b), which makes evident that this modification of the adaptive law limits the capabilities of the MRAC algorithm to improve the transient characteristics and recover desired system performance. Finally, Fig. 20c shows the contribution of the $\mathcal{L}_1$ controller. During the transient phase, the adaptive controller contributes significantly in order to

^{***}In [29], the AP was intentionally (de)tuned to demonstrate performance recovery capabilities of the $\mathcal{L}_1$ augmentation loop.

Fig. 17 MRAC: closed-loop response to biased sinusoidal reference signal in event of left aileron failure and later recovery of full control authority.

stabilize the plant and track the output of the reference system with the available control authority while guaranteeing a smooth recovery. Eventually, when the integrators in the AP readjust and are able to compensate for the failure, the command from the adaptive controller converges to the initial contribution. Figure 20c also shows that the contribution of the $\mathcal{L}_1$ controller is slightly greater for the impaired UAV in order to compensate for the reduction in control authority.

B. Flight-Test Results

1. Nominal Adaptive Loop Performance

Before reproducing Rohrs et al.'s simulations [7,8] in actual flight tests, we first need to verify that the adaptive augmentation loops designed in the HIL simulation environment exhibit satisfactory performance for the nominal plant consisting of the UAV with its AP. To this end, a series of flight tests with the different adaptive

Fig. 18 MRAC with e modification: closed-loop response to biased sinusoidal reference signal in event of left aileron failure and later recovery of full control authority.

Fig. 19 $\mathcal{L}_1$. Closed-loop response to biased sinusoidal reference signal in event of left aileron failure and later recovery of full control authority.

Fig. 20 Lissajous $r_{ad}(r_{cmd})$ for three different adaptive controllers in response to biased sinusoidal reference signal and in the event of left aileron failure and later recovery of full control authority.

controllers were performed in which the closed-loop adaptive system was driven by a biased sinusoidal reference signal of amplitude 7 deg/s at different frequencies.

Figures 21 and 22 present the results obtained in two of these experiments. Both algorithms achieve satisfactory performance, with responses comparable to the ones obtained in HIL simulation. The existence of turbulence makes the quality of the data acquired in flight poorer than in HIL simulation. It is important to mention that, even for the nominal plant without artificially injected unmodeled dynamics, the MRAC parameters slowly diverge. In flight with the $\mathcal{L}_1$ controller, one can observe that at $t = 70$ s, the $\mathcal{L}_1$ adaptive controller generates small oscillations at an approximate frequency of 1 Hz in the control channel for approximately 3 s. Again, these oscillations are due to internal saturation of the $\mathcal{L}_1$ controller. The levels of the projection operator were too tight for the level of turbulence encountered during flight. Nevertheless, stability of the closed-loop adaptive system was not compromised.

2. Second-Order Unmodeled Dynamics

During the flight tests, the same second-order transfer function considered in Sec. IV.A.1, with natural frequency $\omega_n = 1.5$ rad/s

and damping ratio $\zeta = 0.45$, was injected at the output of the actual plant. Using the Lissajous curves, the phase crossover frequency of the real plant with the unmodeled dynamics was determined in real time with data received through telemetry on the ground. Figure 23 shows the Lissajous curve $r(r_{cmd})$ for the system driven by a biased sinusoidal reference signal at $\omega = 1.6$ rad/s; one can see that, similar to the experiments in HIL simulation environment, this Lissajous curve presents a phase shift of approximately -180 deg between the reference signal and the output of the system.

Next, we need to verify that we are able to reproduce in real flight the same instability phenomena observed in HIL simulation. To this end, the system with unmodeled dynamics and the conventional MRAC algorithm is driven by a biased sinusoidal reference signal at the phase crossover frequency. Figure 24 shows the response of the closed-loop adaptive system. One can see that the parameters drift slowly, generating adaptive command signals larger than ± 100 deg/s.

The same experiment was conducted for the MRAC algorithm with the e modification (see Fig. 25). As expected, the parameters do not drift, and the system retains stability during the whole flight. However, similar to HIL results, the performance of the closed-loop

Fig. 21 MRAC: closed-loop nominal response to biased sinusoidal reference signal at $\omega = 0.5$ rad/s.

adaptive system is significantly degraded due to the addition of the damping term in the adaptive law. In this sense, the closed-loop adaptive system is not even able to track the dc component of the reference signal.

Finally, in Fig. 26, we present flight-test results obtained for the $\mathcal{L}_1$ augmentation loop. The system maintains stability during the whole flight and the control signal remains inside reasonable bounds during the experiment. As one would expect, since the frequency of the reference signal is well beyond the bandwidth of the low-pass filter

$C(s)$ in the control law (which was set to 0.6 rad/s), the $\mathcal{L}_1$ controller is not able to achieve desired performance. We note, nevertheless, that the closed-loop system with the $\mathcal{L}_1$ adaptive controller is able to track the bias term of the reference signal, which suggests that the $\mathcal{L}_1$ augmentation loop ensures tracking of the low-frequency content of the reference signal while guaranteeing stability of the system to its high-frequency content. Furthermore, Fig. 27 illustrates the graceful degradation of performance with the $\mathcal{L}_1$ controller as the frequency of the reference signal increases beyond the bandwidth of the low-pass

Fig. 22 $\mathcal{L}_1$: closed-loop nominal response to biased sinusoidal reference signal at $\omega = 0.5$ rad/s.

Fig. 23 Lissajous curve for nominal system with unmodeled dynamics at phase crossover frequency ($\omega = 1.6$ rad/s).

filter. In this set of experiments, the closed-loop system with the second-order unmodeled dynamics and the $\mathcal{L}_1$ controller was driven by biased sinusoidal reference signals at different frequencies. It can be seen that the output of the closed-loop adaptive system is able to track the output of the reference system for reference signals at low frequencies ($\omega = 0.3$ rad/s) and, as the frequency of the reference signal increases, the eccentricity of the Lissajous curve decreases. This graceful degradation in the performance of the system is consistent with the theory of fast and robust adaptation, and it results in a predictable response of the closed-loop adaptive system.

To conclude this section, we present a last set of figures (Figs. 28a and 28b) that demonstrate the recovery capabilities of the $\mathcal{L}_1$ control augmentation loop. The experiment starts with the nominal plant consisting of the UAV and its AP driven by a sinusoidal reference signal at 1.6 rad/s. The second-order unmodeled dynamics are then

a) Lissajous $r(r_m)$

b) Adaptive command to AP

c) MRAC parameters

Fig. 24 MRAC: closed-loop nominal response in presence of second-order unmodeled dynamics to biased sinusoidal reference signal at phase crossover frequency ($\omega = 1.6$ rad/s).

a) Lissajous $r(r_m)$

b) Adaptive command to AP

c) MRAC parameters

Fig. 25 MRAC with e modification: closed-loop nominal response in presence of second-order unmodeled dynamics to biased sinusoidal reference signal at phase crossover frequency ($\omega = 1.6$ rad/s).

a) Lissajous $r(r_m)$

b) Adaptive command to AP

c) $\mathcal{L}_1$ contribution

Fig. 26 $\mathcal{L}_1$: closed-loop nominal response in presence of second-order unmodeled dynamics to biased sinusoidal reference signal at phase crossover frequency ($\omega = 1.6$ rad/s).

injected and, since 1.6 rad/s corresponds to the phase crossover frequency of the system with unmodeled dynamics, the Lissajous figure adopts the typical shape for -180° deg of phase shift. When the MRAC algorithm is engaged, the closed-loop system becomes unstable, as expected, and the control command from the adaptive algorithm starts growing and eventually hits the saturation limit of the AP. Then, at $t = 56$ s, the $\mathcal{L}_1$ controller is engaged and the UAV recovers stability in around 1.5 s, thus confirming the theoretical claims of fast and robust adaptation.

V. Conclusions

This paper presents HIL simulations and flight-test results that evaluate the performance and robustness characteristics of an $\mathcal{L}_1$ control augmentation loop implemented onboard a small UAV. The results in the paper demonstrate that the $\mathcal{L}_1$ controller maintains stability and predictable performance in the presence of unmodeled dynamics and control surface failures. The paper also illustrates the main features of $\mathcal{L}_1$ adaptive control theory and confirms the advantages of $\mathcal{L}_1$ architectures as robust adaptive control architectures with

Fig. 27 $\mathcal{L}_1$: closed-loop response in presence of second-order unmodeled dynamics to a biased sinusoidal reference signal at a–b) $\omega = 0.3$ rad/s, c–d) $\omega = 0.5$ rad/s, e–f) $\omega = 0.7$ rad/s, and g–h) $\omega = 1.0$ rad/s.

Fig. 28 MRAC and $\mathcal{L}_1$ algorithms in flight reproducing Rohrs et al.'s setup [7,8].

the potential of facilitating the transition of adaptive control into advanced flight control systems.

From a historical perspective, this investigation was motivated by the Rohrs counterexample, a benchmark problem presented in the early 1980s to show the deficiencies of adaptive controllers developed at that time and for which the main criticisms regarding performance and robustness of conventional MRAC architectures had remained unresolved. The results presented in the paper demonstrate that the flight-test setup qualitatively replicates the results obtained by Rohrs et al. [7,8]. In fact, the study reproduces the lack of predictability of conventional MRAC algorithms and shows that this class of adaptive algorithms may unexpectedly lose stability in the presence of unmodeled dynamics. In addition, it has been shown that the damping-type modifications in the adaptive laws proposed in the literature retain stability of the closed-loop adaptive system but introduce a significant degradation in the tracking capabilities of the closed-loop system.

$\mathcal{L}_1$ adaptive control overcomes these deficiencies by decoupling adaptation from robustness, which is in contrast with the current prevailing conventional wisdom that slow adaptation is needed so that robustness would not be compromised. In $\mathcal{L}_1$ adaptive control architectures, robustness of the closed-loop adaptive system is ensured by the low-pass filter inserted in the control law, while fast adaptation results in a predictable and consistent closed-loop performance. In this sense, $\mathcal{L}_1$ adaptive control solves the problem that Rohrs et al. [7,8] posed regarding performance and robustness of conventional adaptive architectures.

Appendix: Projection-Type Operator in Adaptation Laws

Let θ be an unknown parameter, and assume that it belongs to the convex compact set Θ ; that is, $\theta \in \Theta$. Then, the $\text{Proj}: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^n$ operator is defined as

$$\text{Proj}(\hat{\theta}, x) = \begin{cases} x & \text{if } h(\hat{\theta}) < 0 \\ x & \text{if } h(\hat{\theta}) > 0, \nabla h^\top x \leq 0 \\ x - \frac{\nabla h \nabla h^\top y h(\hat{\theta})}{\|\nabla h\|^2} & \text{if } h(\hat{\theta}) > 0, \nabla h^\top x > 0 \end{cases}$$

where

$$h(\hat{\theta}) = \frac{\hat{\theta}^\top \hat{\theta} - \theta_{\max}^2}{\epsilon_\theta \theta_{\max}}$$

with $\theta_{\max} \in \mathbb{R}^+$ being the norm bound imposed on $\hat{\theta}$, and $\epsilon_\theta \in \mathbb{R}^+$ being the convergence tolerance of the bound. By appropriately choosing $\theta_{\max}$ and ϵ_θ , the Proj operator ensures that $\hat{\theta}(t) \in \Theta$ for all $t \geq 0$ [35]. The following property of the projection operator is used in the corresponding Lyapunov analysis.

Lemma [35]: Given $x, \hat{\theta} \in \mathbb{R}^n$, we have

$$(\hat{\theta} - \theta)^\top [\text{Proj}(\hat{\theta}, x) - x] \leq 0$$

where θ is the true value of $\hat{\theta}$.

Acknowledgment

This work was sponsored in part by NASA grants NNX08AB97A, NNX08AC81A, and NNL08AA12I.

References

- [1] Anderson, B. D. O., "Failures of Adaptive Control Theory and their Resolution," *Communications in Information and Systems*, Vol. 5, No. 1, 2005, pp. 1–20.
- [2] Anderson, B. D. O., and Dehghani, A., "Challenges of Adaptive Control: Past, Permanent and Future," *Annual Reviews in Control*, Vol. 32, No. 2, 2008, pp. 123–135. doi:10.1016/j.arcontrol.2008.06.001
- [3] Wise, K. A., Lavretsky, E., and Hovakimyan, N., "Adaptive Control in Flight: Theory, Application, and Open Problems," *American Control Conference*, Minneapolis, MN, IEEE Publ., Piscataway, NJ, June 2006, pp. 5966–5971.
- [4] Jacklin, S. A., Lowry, M. R., Schumann, J. M., Gupta, P. P., Bosworth, J. T., Zavala, E., Kelly, J. W., Hayhurst, K. J., Belcastro, C. M., and Belcastro, C. M., "Verification, Validation, and Certification Challenges for Adaptive Flight-Critical Control System Software," *AIAA Guidance, Navigation and Control Conference*, Providence, RI, AIAA Paper 2004-5258, Aug. 2004.
- [5] Jacklin, S. A., "Closing Certification Gaps in Adaptive Flight Control Software," *AIAA Guidance, Navigation and Control Conference*, Honolulu, HI, AIAA Paper 2008-6988, Aug. 2008.
- [6] Egardt, B., "Stability of Model Reference Adaptive and Self-Tuning Regulators," Ph.D. Thesis, Lund Inst. of Technology, Lund, Sweden, Dec. 1978.
- [7] Rohrs, C. E., Valavani, L. S., Athans, M., and Stein, G., "Robustness of Adaptive Control Algorithms in the Presence of Unmodeled Dynamics," *IEEE Conference on Decision and Control*, Orlando, FL, Vol. 1, IEEE Publ., Piscataway, NJ, Dec. 1982, pp. 3–11.
- [8] Rohrs, C. E., Valavani, L. S., Athans, M., and Stein, G., "Robustness of Continuous-Time Adaptive Control Algorithms in the Presence of Unmodeled Dynamics," *IEEE Transactions on Automatic Control*, Vol. 30, No. 9, Sept. 1985, pp. 881–889. doi:10.1109/TAC.1985.1104070
- [9] Georgiou, T. T., and Smith, M. C., "Robustness Analysis of Nonlinear Feedback Systems: An Input–Output Approach," *IEEE Transactions on Automatic Control*, Vol. 42, No. 9, Sept. 1997, pp. 1200–1221. doi:10.1109/9.623082
- [10] Anderson, B. D. O., "Adaptive Systems, Lack of Persistency of Excitation and Bursting Phenomena," *Automatica*, Vol. 21, No. 3, May 1985, pp. 247–258. doi:10.1016/0005-1098(85)90058-5
- [11] Mareels, I. M. Y., and Bitmead, R. R., "Non-linear Dynamics in Adaptive Control: Chaotic and Periodic Stabilization," *Automatica*, Vol. 22, No. 6, Nov. 1986, pp. 641–655. doi:10.1016/0005-1098(86)90003-8

- [12] Aloneftis, A., *Stochastic Adaptive Control: Results and Simulations*, Vol. 98 of Lecture Notes in Control and Information Sciences, Springer-Verlag, New York, 1987, pp. 39–58.
- [13] Kosut, R. L., Mareels, I. M. Y., Anderson, B. D. O., Bitmead, R. R., and Johnson, C. R., Jr., “Transient Analysis of Adaptive Control,” *IFAC World Congress*, Munich, Elsevier, Oxford, England, U.K., July 1987, pp. 121–126.
- [14] Mareels, I. M. Y., and Bitmead, R. R., “Non-linear Dynamics in Adaptive Control: Chaotic and Periodic Stabilization: II. Analysis,” *Automatica*, Vol. 24, No. 4, July 1988, pp. 485–497. doi:10.1016/0005-1098(88)90093-3
- [15] Zang, Z., and Bitmead, R. R., “Transient Bounds for Adaptive Control Systems,” *IEEE Conference on Decision and Control*, Vol. 5, Honolulu, HI, IEEE Publ., Piscataway, NJ, Dec. 1990, pp. 2724–2729.
- [16] Åström, K. J., “Analysis of Rohrs Counterexamples to Adaptive Control,” *IEEE Conference on Decision and Control*, San Antonio, TX, IEEE Publ., Piscataway, NJ, Dec. 1983, pp. 982–987.
- [17] Åström, K. J., “Interactions between Excitation and Unmodeled Dynamics in Adaptive Control,” *IEEE Conference on Decision and Control*, Las Vegas, NV, IEEE Publ., Piscataway, NJ, Dec. 1984, pp. 1276–1281.
- [18] Ioannou, P. A., and Kokotović, P. V., “An Asymptotic Error Analysis of Identifiers and Adaptive Observers in the Presence of Parasitics,” *IEEE Transactions on Automatic Control*, Vol. 27, No. 4, Aug. 1982, pp. 921–927. doi:10.1109/TAC.1982.1103028
- [19] Peterson, B. B., and Narendra, K. S., “Bounded Error Adaptive Control,” *IEEE Transactions on Automatic Control*, Vol. 27, No. 6, Dec. 1982, pp. 1161–1168. doi:10.1109/TAC.1982.1103112
- [20] Kresselmeier, G., and Narendra, K. S., “Stable Model Reference Adaptive Control in the Presence of Bounded Disturbances,” *IEEE Transactions on Automatic Control*, Vol. 27, No. 6, Dec. 1982, pp. 1169–1175. doi:10.1109/TAC.1982.1103093
- [21] Ioannou, P. A., and Kokotović, P. V., *Adaptive Systems with Reduced Models*, Springer-Verlag, Secaunus, NJ, 1983, pp. 81–89.
- [22] Ioannou, P. A., and Kokotović, P. V., “Robust Redesign of Adaptive Control,” *IEEE Transactions on Automatic Control*, Vol. 29, No. 3, March 1984, pp. 202–211. doi:10.1109/TAC.1984.1103490
- [23] Narendra, K. S., and Annaswamy, A. M., “A New Adaptive Law for Robust Adaptation Without Persistent Excitation,” *IEEE Transactions on Automatic Control*, Vol. 32, No. 2, February 1987, pp. 134–145. doi:10.1109/TAC.1987.1104543
- [24] Ortega, R., and Tang, Y., “Robustness of Adaptive Controllers: A Survey,” *Automatica*, Vol. 25, No. 5, Sept. 1989, pp. 651–677. doi:10.1016/0005-1098(89)90023-X
- [25] Hovakimyan, N., and Cao, C., *$\mathcal{L}_1$ Adaptive Control Theory*, Soc. for Industrial and Applied Mathematics, Philadelphia, 2010.
- [26] Dobrokhodov, V., Yakimenko, O., Jones, K. D., Kaminer, I., Bourakov, E., Kitsios, I., and Lizarraga, M., “New Generation of Rapid Flight Test Prototyping System for Small Unmanned Air Vehicles,” AIAA Modelling and Simulation Technologies Conference, Hilton Head Island, SC, AIAA Paper 2007-6567, Aug. 2007.
- [27] Cao, C., and Hovakimyan, N., “ $\mathcal{L}_1$ Adaptive Output Feedback Controller for Systems of Unknown Dimension,” *IEEE Transactions on Automatic Control*, Vol. 53, No. 3, April 2008, pp. 815–821. doi:10.1109/TAC.2008.919550
- [28] Cao, C., and Hovakimyan, N., “ $\mathcal{L}_1$ Adaptive Output-Feedback Controller for Non-Strictly-Positive-Real Reference Systems: Missile Longitudinal Autopilot Design,” *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 3, May–June 2009, pp. 717–726. doi:10.2514/1.40877
- [29] Dobrokhodov, V., Kitsios, I., Kaminer, I., Jones, K. D., Xargay, E., Hovakimyan, N., Cao, C., Lizarraga, M. I., and Gregory, I. M., “Flight Validation of Metrics Driven $\mathcal{L}_1$ Adaptive Control,” AIAA Guidance, Navigation and Control Conference, Honolulu, HI, AIAA Paper 2008-6987, Aug. 2008.
- [30] Cundy, H. M., and Rollett, A. P., “Lissajous’s Figures,” *Mathematical Models*, 3rd ed., Sec. 5.5.3, Tarquin Publ., Stradbroke, England, U.K., 1989, pp. 242–244.
- [31] Dobrokhodov, V., Kitsios, I., Kaminer, I., Jones, K. D., Xargay, E., Hovakimyan, N., Cao, C., Lizarraga, M. I., and Gregory, I. M., “Preliminary Results of Development, System Integration and Flight Validation of a Metrics Driven $\mathcal{L}_1$ Adaptive Control,” AIAA Infotech@Aerospace, Seattle, WA, AIAA Paper 2009-2053, April 2009.
- [32] Michini, B., and How, J., “ $\mathcal{L}_1$ Adaptive Control for Indoor Autonomous Vehicles: Design Process and Flight Testing,” AIAA Guidance, Navigation and Control Conference, Chicago, IL, AIAA Paper 2009-5754, Aug. 2009.
- [33] Narendra, K. S., and Annaswamy, A. M., *Stable Adaptive Systems, Information and System Sciences*, Prentice-Hall, Englewood Cliffs, NJ, 1989.
- [34] Ioannou, P. A., and Sun, J., *Robust Adaptive Control*, Prentice-Hall, Upper Saddle River, NJ, 1996.
- [35] Pomet, J.-B., and Praly, L., “Adaptive Nonlinear Regulation: Estimation from the Lyapunov Equation,” *IEEE Transactions on Automatic Control*, Vol. 37, No. 6, June 1992, pp. 729–740. doi:10.1109/9.256328