Adaptive Flight Control Experiments using RAVEN

Jonathan P. How, Justin Teo, and Buddy Michini

Aerospace Controls Laboratory
Massachusetts Institute of Technology, Cambridge, MA, 02139

***Abstract*— This paper discusses a new flight test facility at MIT known as the Real-time indoor Autonomous Vehicle test ENvironment (RAVEN) which enables the exploration of aggressive aerobatic flight with relatively low risk. A key feature of this work is that the control algorithms are often the limiting factor because the flight dynamics are highly nonlinear and not well known, and they can change with the flight mode and over time. This paper discusses the RAVEN facility and presents some initial experimental results using fixed wing UAVs. These results demonstrate some of the key advantages of using RAVEN; illustrate the types of performance limitations typically faced when executing aggressive aerobatic maneuvers; and motivate the current research on adaptive flight control.**

I. Introduction

Unmanned Aerial Vehicles (UAVs) are becoming vital warfare and homeland security platforms because they significantly reduce costs and the risk to human life while amplifying warfighter and first-responder capabilities [1]. Given these recent application successes, our work is motivated by future mission scenarios that involve flying micro (and/or nano) air vehicles (called MAVs and NAVs) in highly constrained environments. Typical examples of interest include MAV flight in an urban canyon and an NAV transitioning from outdoor flight to indoors. With these goals in mind, the objectives of this work are to develop and validate flight control concepts for aggressive (aerobatic) maneuvers, and, in particular, to identify the sensor suites needed, and the likely limits of achievable performance. Our approach includes flight testing in a unique facility called the Real-time indoor Autonomous Vehicle test ENvironment (RAVEN) at MIT's Aerospace Controls Laboratory [2], [3].

References [4], [5] discuss several of the multi-UAV platforms that have been developed by numerous research groups for both outdoor and indoor use. Those papers also highlight some of the key limitations and benefits of these testbeds. In particular, in addition to weather and daylight constraints, external UAVs typically require a large safety and support team, which makes testing logistically difficult and expensive. The conclusions of this analysis are that, while there are some limitations to indoor flight using RAVEN (e.g., limited space, some communication noise), we believe that there are significant advantages to this approach, namely: information-rich, controlled environment, logistics, safety, and cost scalability. Furthermore, since the system software is designed to autonomously manage the navigation and tasking of the air vehicles during operations, researchers can focus on the algorithms associated with the research project (e.g., flight control or team coordination). These properties greatly enhance the utility of the testbed, making it an extremely effective **rapid prototyping** and **validation** environment for aircraft flight control.

Given these capabilities, we have been investigating flight control algorithms of MAVs and NAVs for operations in urban environments that require high maneuverability and possibly non-traditional vehicle configurations. The control design problem to achieve this agile flight is very complex and there are numerous technical issues, such as: (i) the aerodynamic, vehicle, and actuator modeling is very difficult and often results in a high degree of uncertainty; (ii) must handle multiple flight modes and vehicle configurations; and (iii) vehicle sensing in GPS-denied environments is very challenging. These issues motivate further work on a combination of adaptive, nonlinear, and hybrid control, and they also help motivate the development of a robust flight facility such as RAVEN that enables rapid prototyping of these control techniques.

II. Comparison between Proportional-Integral Control and Direct Adaptive Approximate Dynamic Inversion

The authors of [6], [7] laid the foundation for a method of Approximate Dynamic Inversion (ADI) for a class of minimum phase, nonaffine-in-control systems, assuming known system dynamics. The method is founded on the time-scale separation principle from singular perturbation theory [8], where the control is defined as a solution of "fast" dynamics. The ADI control law as originally formulated depends on the nonlinear function that describes the system. When this function is known, implementation of the ADI controller is straightforward. When this function is unknown, one plausible way would be to estimate this function and construct an analogous ADI control law based on that estimate [9]–[11]. A Direct Adaptive approach is proposed in [12] to address the same issue.

Professor in the Dept. of Aeronautics and Astronautics, MIT, 77 Mass. Ave., Cambridge, MA 02139, jhow@mit.edu

Ph.D. Candidate, Department of Aeronautics and Astronautics, csteo@mit.edu

S.M. Candidate, Department of Aeronautics and Astronautics, bmich@mit.edu

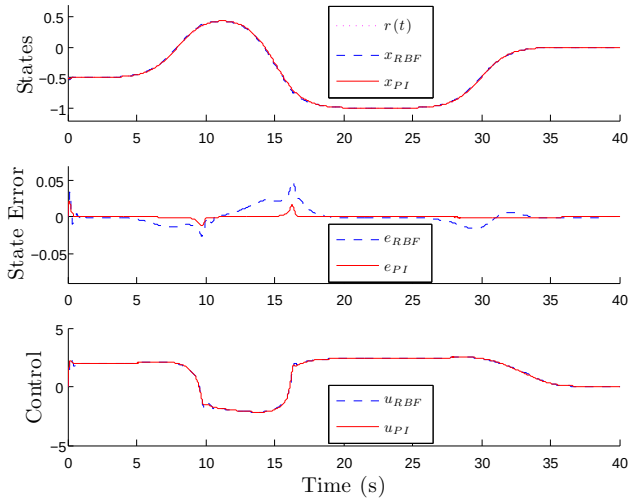

Fig. 1. Simulation Results for $r(t) = r_a(t)$

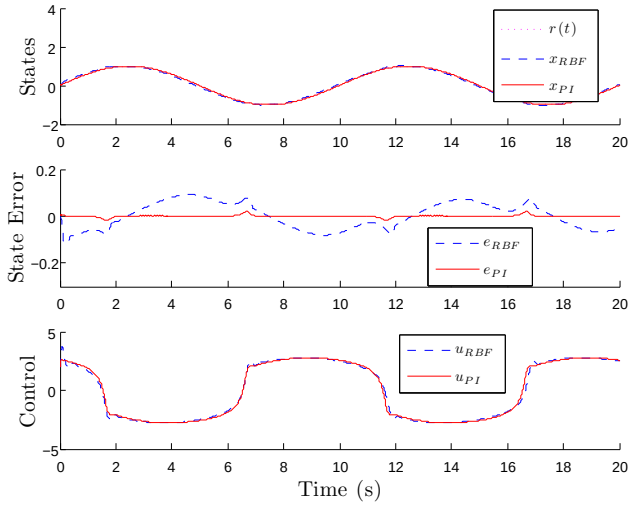

Fig. 2. Simulation Results for $r(t) = r_b(t)$

Previous work by the authors in [13] showed that the ADI control law is actually equivalent to a Proportional-Integral (PI) controller, which can be implemented exactly with full state feedback and only knowledge of the sign of the control effectiveness. To illustrate the main idea, consider a first order single input, minimum phase nonaffine-in-control system described by

$$\dot{x}(t) = f(x(t), u(t)).$$

Define the *stable* linear reference model

$$\dot{x}_r(t) = -a_r x_r(t) + b_r r(t)$$

with $a_r > 0$, and tracking error $e(t) = x(t) - x_r(t)$. Here, $r(t)$ is the reference signal that drives the reference model. It is desired for the states of the system, $x(t)$, to track the states of the reference model, $x_r(t)$. The ADI control law is then given by

$$\epsilon \dot{u}(t) = -\operatorname{sign}\left(\frac{\partial f}{\partial u}\right)\left(f\left(x(t), u(t)\right) + a_r x(t) - b_r r(t)\right). \tag{1}$$

Recognizing that $f(x(t), u(t)) = \dot{x}(t)$, (1) can be rewritten as

$$\epsilon \dot{u}(t) = -\operatorname{sign}\left(\frac{\partial f}{\partial u}\right)\left(\dot{x}(t) + a_r x(t) - b_r r(t)\right),$$

which can be implemented as

$$u(t) = -\frac{1}{\epsilon}\operatorname{sign}\left(\frac{\partial f}{\partial u}\right)\left(x(t) + \int_0^t a_r x(\tau) - b_r r(\tau)\,\mathrm{d}\,\tau\right). \tag{2}$$

The control law (2) is effectively a PI controller that realizes the ADI controller without requiring any approximations of $f(x(t), u(t))$. Observe that the only knowledge required of the system is the sign of the control effectiveness, $\operatorname{sign}\left(\frac{\partial f}{\partial u}\right) \in \{-1, 1\}$.

Simulation results in [13] comparing the Indirect Adaptive controller of [11] and the equivalent PI controller shows that the PI controller achieves/exceeds the tracking performance of the Indirect Adaptive controller.

In this section, we extract the equivalent PI controller from the Direct Adaptive controller of [12], and compare their tracking performance on the example system in [12].

A. Simulation Results

In [12], the example nonaffine-in-control system is described by

$$\dot{x}(t) = 0.5x(t) + \tanh(u(t) + 3)\tanh(u(t) - 3) + 0.01u(t).$$

The (partial) set of equations describing the Direct Adaptive control law of [12] are

$$u(t) = 2\nu(t) \tag{3a}$$

$$\nu(t) = -10e(t) + \dot{r}(t) - \nu_{ad}(t) \tag{3b}$$

$$0.05\dot{\nu}_{ad}(t) = \hat{W}^{\mathrm{T}}\Phi(x(t), u(t)) + 10e(t) - \dot{r}(t) \tag{3c}$$

where, $e(t) = x(t) - r(t)$, both $r(t)$ and $\dot{r}(t)$ as assumed to be available for feedback, and $\hat{W}^{\mathrm{T}}\Phi(x(t), u(t))$ is a Radial Basis Function (RBF) approximator for $f(x(t), u(t))$. Additional equations describing the complete Direct Adaptive controller with regards to training the RBF Neural Network are not presented here for simplicity, and to avoid clutter.

To extract the equivalent PI controller, observe that $\hat{W}^{\mathrm{T}}\Phi(x(t), u(t))$ is used to estimate $f(x(t), u(t)) = \dot{x}(t)$. Under a perfect estimation scenario, (3c) can be written as

$$\begin{aligned} 0.05\dot{\nu}_{ad}(t) &= f(x(t), u(t)) + 10e(t) - \dot{r}(t) \\ &= \dot{x}(t) + 10e(t) - \dot{r}(t) \\ &= \dot{e}(t) + 10e(t), \end{aligned}$$

and implemented as

$$\nu_{ad}(t) = 20e(t) + 200\int_0^t e(\tau)\,\mathrm{d}\,\tau.$$

Collecting (3) with the above replacement for $\nu_{ad}(t)$, the equivalent PI controller is then given by

$$u(t) = 2\dot{r}(t) - 60e(t) - 400\int_0^t e(\tau)\,\mathrm{d}\,\tau.$$

Observe that no knowledge of $f(x(t), u(t))$ is required in the above and the perfect estimation assumption used to extract the PI controller is valid.

Fig. 1 shows the simulation results for $r(t) = r_a(t)$, with

$$r_a(t) = -\frac{1}{1+e^{t-8}} + \frac{1.5}{1+e^{t-15}} + \frac{1}{1+e^{t-30}}.$$

Fig. 2 shows the simulation results for $r(t) = r_b(t)$, with

$$r_b(t) = \sin(0.2\pi t).$$

In each of these, the top diagram shows the reference, $r(t)$, together with the state of the system, $x(t)$. The middle diagram shows the state error, $e(t)$, and the bottom diagram shows the control, $u(t)$. Signals associated with the Direct Adaptive and PI controller are denoted with subscripts "RBF" and "PI" respectively. Observe that the PI controller achieves/exceeds the tracking performance of the Direct Adaptive controller with a significantly simpler implementation.

III. RAVEN

Fig. 3 presents the control architecture used for all of the flight experiments given in the paper [5]. A key feature of RAVEN is the Vicon MX motion capture system [14], [15] that can accurately track all vehicles in the room in real-time. With lightweight reflective balls attached to each vehicle's structure, the Vicon motion capture system can measure the vehicle's position and attitude information at rates up to 120 Hz, with approximately a 10 ms delay, and sub-mm accuracy [14]. The bright LED rings of five of the eighteen Vicon cameras in RAVEN are visible in Fig. 4. Just as GPS spurred the development of large-scale UAVs, we expect this new sensing capability to have a significant impact on 3D indoor flight.

The RAVEN facility follows the same design philosophy used in the previous MIT ACL testbeds [16] in that the perception and planning computation is done off-board. Fig. 3 shows that RAVEN actually takes this philosophy one step further by also computing vehicle flight control commands off-board. These commands are computed using ground-based computers at rates that exceed 50 Hz and sent to the vehicles via standard Radio Control (R/C) transmitters. An important feature of this setup is that small, inexpensive, essentially unmodified, radio-controlled vehicles can be used. This enables researchers to avoid being overly conservative during flight testing.

a) Flight Hardware: Numerous aircraft have been used in the recent flight tests in RAVEN. To create a light, acrobatic airframe that can be easily flown indoors, a three-wing tailsitter has been designed and built. As shown

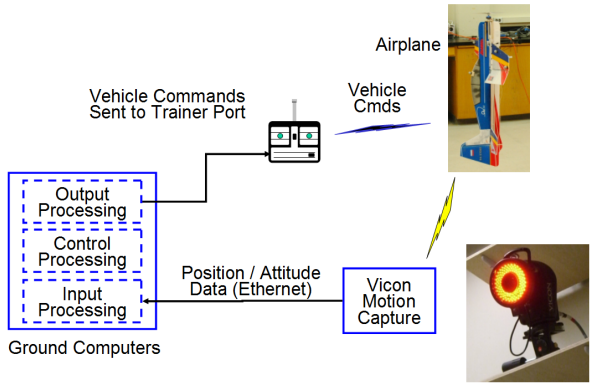

Fig. 3. VICON motion capture camera. Lens located at center of the LED-array.

Fig. 4. RAVEN flight space.

in Fig. 5, a propeller is mounted above three triangular, symmetrical wings with large flaps at the bottom of each. The onboard electronics are commercial off-the-shelf RC aircraft components and Vicon dots are added to enable position and attitude measurements. The vehicle is capable of vertical takeoff/landing as well as both hover and horizontal flight.

Given this ability to easily transition between hover and horizontal flight modes, the ultimate control goal for the vehicle is a single, unified controller for all flight regimes. To foster this, a quaternion-based attitude reference system has been adapted and implemented. A quaternion representing the desired attitude is generated by the path planner and then compared to the quaternion representing the aircraft's actual attitude. An error quaternion is calculated, the vector components of which represent generic roll, pitch, and yaw angular errors in the body frame of the aircraft. Inner-loop controllers then use these errors to generate roll, pitch, and yaw commands. An advantage of this inner-loop control structure is that it is not airframe-specific. In the case of the three-wing, the generic commands are simply translated into the appropriate actuation sets.

b) Preliminary Adaptive Control Results: An $\mathcal{L}_1$-adaptive controller has been implemented using the techniques discussed in [17], [18] to control the yaw axis of

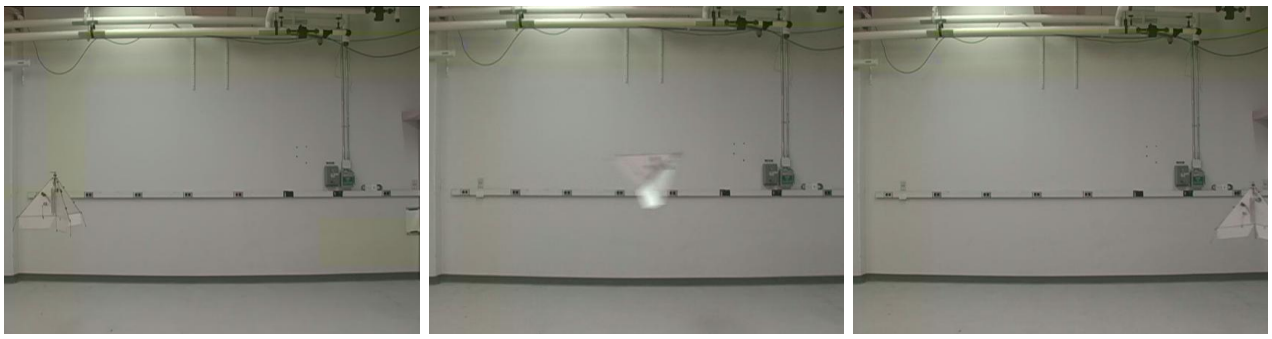

(a) (b) (c)

Fig. 5. The three-wing vehicle hovering at the start position (a), transition to near-horizontal flight (b), hovering at final position (c).

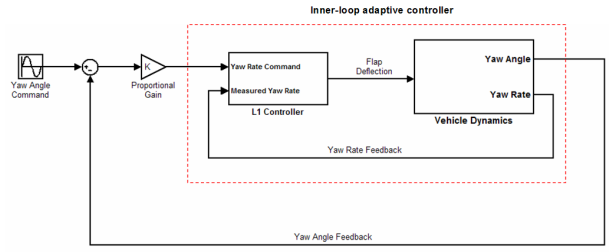

Fig. 6. Configuration for $\mathcal{L}_1$-adaptive yaw axis control.

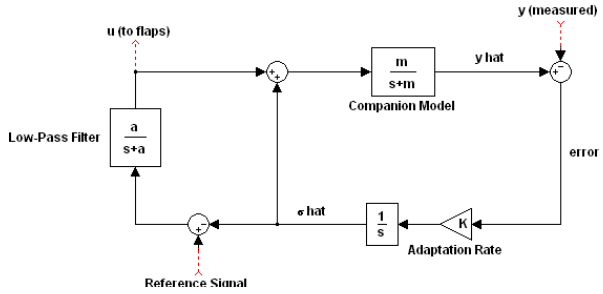

Fig. 7. $\mathcal{L}_1$ output feedback block diagram.

the aircraft (the propeller axis). As shown in Fig. 6, the $\mathcal{L}_1$ controller only regulates the yaw rate as an inner-loop controller. An outer-loop linear proportional controller then uses the yaw angular error to feed a rate command to the $\mathcal{L}_1$ controller. The motivation for this design is to provide a layer of separation between the linear controller and the vehicle dynamics. Assuming the $\mathcal{L}_1$-adaptive controller performs as intended, it will force the vehicle dynamics to be similar to some chosen (stable, minimum phase) reference model, which allows the linear controller to be designed using these known dynamics. Ideally, this single linear controller could then be used in numerous flight modes.

Fig. 7 shows the structure of the inner-loop adaptive controller. It is a first-order version of the $\mathcal{L}_1$ controllers formulated in [17], [18] intended by the authors to be used as an output-feedback controller for systems of unknown dimension, so it is well suited for this inner-loop application. It is also computationally simple to implement, consisting of only two first-order lag filters (the low-pass filter and companion model) and an output-limited integrator. It should be noted that the companion model is not equivalent to the reference model in typical model reference adaptive systems since the high-frequency component of the adaptive control input is fed in along with the reference signal. This implies that the output of the companion model is affected by both the adaptation rate and low-pass filter. As a result, the companion model is not an independent performance reference for the system as in conventional MRAC. Instead, selection of the companion model depends on the choice of adaptive rate and low-pass filter.

Fig. 8 compares a baseline linear yaw controller and the $\mathcal{L}_1$-adaptive yaw controller for nominal flight conditions under a sinusoidal yaw angle command. Fig. 8(a) shows the yaw angle tracking. The controllers demonstrate similar performance in the nominal case. It should be noted that the controllers were not tuned to match performance, but rather to yield acceptable tracking. In Fig. 8(b), the dashed blue line is the reference signal (an angular rate command from the outer loop), while the green line is the companion model output (the signal to be tracked). The red line is the measured angular rate, which matches the green almost exactly implying that the inner-loop $\mathcal{L}_1$ controller tracks the companion model very closely. To investigate the ability of the $\mathcal{L}_1$ controller to adapt to variations in plant dynamics, the actuator effectiveness was reduced to 50% of its original value (i.e. the servo commands were multiplied by 0.5). Fig. 9 shows that the linear controller performance is quite poor, whereas the adaptive controller maintains performance similar to the nominal case.

IV. Summary

This paper presented some recent analysis of the direct adaptive approximate dynamic inversion algorithm and illustrated that RAVEN provides a great facility to rapid prototype adaptive flight controllers for challenging vehicles and mission scenarios.

Acknowledgments

The authors would like to thank Frantisek Sobolic, Adrian Frank, Spencer Ahrens, Brett Bethke, Eli Cohen,

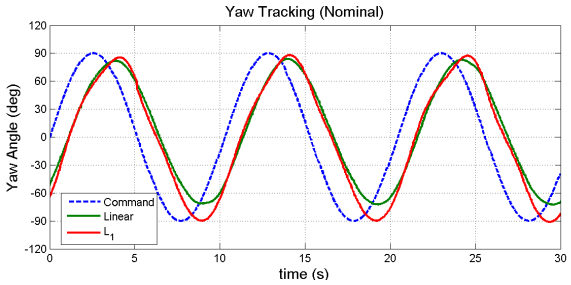

(a)

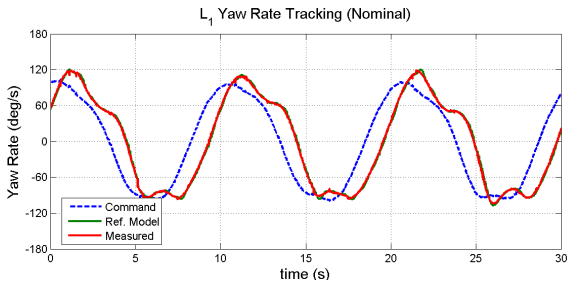

(b)

Fig. 8. Comparison of linear and $\mathcal{L}_1$ yaw controllers (a) and $\mathcal{L}_1$ yaw rate tracking (b), nominal conditions.

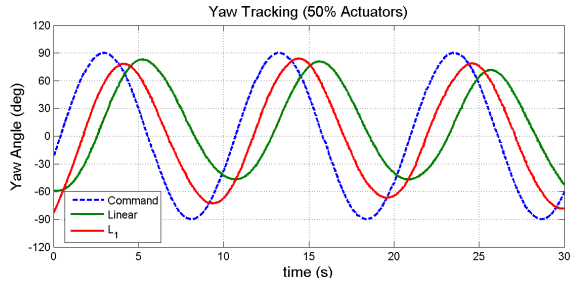

(a)

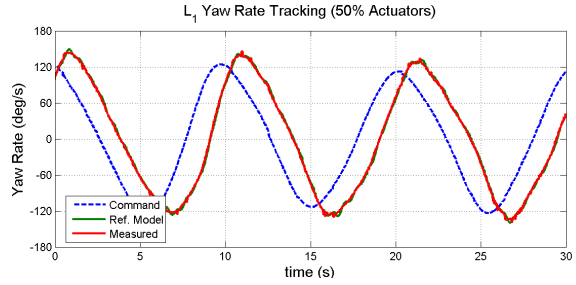

(b)

Fig. 9. Comparison of linear and $\mathcal{L}_1$ yaw controllers (a) and $\mathcal{L}_1$ yaw rate tracking (b), 50% actuator effectiveness.

Daniel Dale, Brandon Luders, and Adam Woodworth for their invaluable assistance. Research supported in part by AFOSR grant FA9550-04-1-0458 and by the Boeing Company, Phantom Works, Seattle.

REFERENCES

[1] Office of the Secretary of Defense, "Unmanned Aircraft Systems Roadmap, `http://www.acq.osd.mil/usd/UnmannedSystemsRoadmap.2007-2032.pdf`," tech. rep., OSD, 2007.

[2] M. Valenti, B. Bethke, D. Dale, A. Frank, J. McGrew, S. Ahrens, J. P. How, and J. Vian, "The MIT Indoor Multi-Vehicle Flight Testbed," in *Proceedings of the IEEE International Conference on Robotics and Automation*, (Rome, Italy), April 2007. Video Submission.

[3] M. Valenti, B. Bethke, J. How, and J. Vian, "Embedding health management into mission tasking for UAV teams," in *Proceedings of the American Control Conference*, (New York, NY), June 2007.

[4] J. P. How, "Multi-vehicle flight experiments: Recent results and future directions," in *Symposium on Platform Innovations and System Integration for Unmanned Air, Land and Sea Vehicles*, vol. AVT-146, (Florence, Italy), NATO, May 2007.

[5] J. P. How, B. Bethke, A. Frank, D. Dale, and J. Vian, "Real-time indoor autonomous vehicle test environment," *Control Systems Magazine (*to appear*)*, 2008.

[6] N. Hovakimyan, E. Lavretsky, and A. Sasane, "Dynamic inversion for nonaffine-in-control systems via time-scale separation. part i," *Journal of Dynamical and Control Systems*, vol. 13, pp. 451 – 465, October 2007.

[7] N. Hovakimyan, E. Lavretsky, and A. J. Sasane, "Dynamic inversion for nonaffine-in-control systems via time-scale separation: Part i," in *Proceedings of the American Control Conference*, (Portland, OR), pp. 3542 – 3547, June 2005.

[8] H. K. Khalil, *Nonlinear Systems*. Prentice Hall, 3rd ed., 2002.

[9] E. Lavretsky and N. Hovakimyan, "Adaptive dynamic inversion for nonaffine-in-control uncertain systems via time-scale separation. part ii," *Journal of Dynamical and Control Systems*, vol. 14, pp. 33 – 41, January 2008.

[10] E. Lavretsky and N. Hovakimyan, "Adaptive dynamic inversion for nonaffine-in-control systems via time-scale separation: Part ii," in *Proceedings of the American Control Conference*, (Portland, OR), pp. 3548 – 3553, June 2005.

[11] N. Hovakimyan, E. Lavretsky, and C. Cao, "Adaptive dynamic inversion via time-scale separation," in *Proceedings of the 45th IEEE Conference on Decision & Control*, (San Diego, CA), pp. 1075 – 1080, December 2006.

[12] E. Lavretsky and N. Hovakimyan, "Adaptive compensation of control dependent modeling uncertainties using time-scale separation," in *Proceedings of the 44th IEEE Conference on Decision and Control and the European Control Conference*, (Seville, Spain), pp. 2230 – 2235, December 2005.

[13] J. Teo and J. P. How, "Equivalence between approximate dynamic inversion and proportional-integral control," in *47th IEEE Conference on Decision and Control*, (Cancun, Mexico), December 2008. submitted for publication.

[14] M. Valenti, B. Bethke, G. Fiore, J. How, and E. Feron, "Indoor multi-vehicle flight testbed for fault detection, isolation, and recovery," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference and Exhibit*, (Keystone, CO), August 2006. AIAA-2006-6200.

[15] Vicon, "Vicon MX Systems." Available at `http://www.vicon.com/products/viconmx.html`, July 2006.

[16] E. King, Y. Kuwata, and J. P. How, "Experimental demonstration of coordinated control for multi-vehicle teams," *International Journal of Systems Science*, vol. 37, pp. 385–398, May 2006.

[17] C. Cao and N. Hovakimyan, "Design and analysis of a novel l1 adaptive controller, part i: Control signal and asymptotic stability," in *American Control Conference*, (Minneapolis, MN), pp. 3397 – 3402, IEEE, June 2006.

[18] C. Cao and N. Hovakimyan, "Design and analysis of a novel l1 adaptive controller, part ii: Guaranteed transient performance," in *American Control Conference*, (Minneapolis, MN), pp. 3403 – 3408, IEEE, June 2006.