Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 54-12 (2021) 14–19

A course on adaptive and autonomous aerospace systems

Davide Invernizzi * **Marco Lovera** *

* *Dipartimento di Scienze e Tecnologie Aerospaziali, Politecnico di Milano, Via La Masa 34, 20156, Milano, Italy (e-mail: {davide.invernizzi,marco.lovera}@polimi.it).*

Abstract: In this paper we present the one-semester course "Adaptive and Autonomous aerospace Systems" to be taught to master students in aerospace engineering at Politecnico di Milano starting fall 2021. The paper discusses the course motivation and objectives, how it fits within the current study program in aerospace engineering, the course syllabus and organization and the learning assessment and outcomes. While being a fundamental course in character, covering the main theoretical aspects of adaptive and autonomous control, laboratory activities will complement the students learning by showing realistic implementations of the algorithms taught in class.

1. INTRODUCTION

Automatic control systems play an increasingly important role in aerospace engineering, both in view of the higher level of automation and autonomy expected from flight vehicles and of the recent emergence of unmanned aerial vehicles (UAVs). In particular, control systems design problems in aerospace pose significant challenges because of their intrinsically multivariable, nonlinear nature, often associated with large model uncertainty and unstable dynamics. These are the main reasons why advanced methods for analysis and synthesis must be adopted in aerospace applications. The use of robust design and analysis methods is becoming the standard industrial practice in the aerospace field, leveraging efficient and mature algorithms implemented in commercial software such as the Mathworks' Robust Control Toolbox. Modern robust control has gained a significant interest in aircraft applications also in view of recent developments, which allow robust tuning of consolidated control architectures against multiple design requirements (Apkarian et al., 2014).

When considering the development of high performance autopilots for highly maneuverable vehicles or the presence of unpredictable flight conditions, nonlinear phenomena and time-varying effects significantly alter the dynamical models, making linear design methods inadequate. Adaptive control has emerged as one of the design approaches to solving flight control problems in these challenging conditions (Lavretsky and Wise, 2013; Johnson and Kannan, 2005; Dydek et al., 2013; Mallikarjunan et al., 2012). Current trends in aerospace involve the design of adaptive controllers capable of restoring a desired closed-loop behavior regardless of structured parametric uncertainties, while guaranteeing robustness to unmodeled but bounded disturbances.

At the same time, autonomy, which can be defined, broadly speaking, as the ability of a system to achieve goals while operating independently of external controls, is needed in the presence of communication constraints (*e.g.*, delays, bandwidth and communication windows) and when decisions are time-critical or when they are better made using onboard data compared to downloaded data. In the aerospace sector (Beard et al., 2005), the development of autonomous systems is steadily growing, with applications ranging from air traffic management to autonomous UAVs exploration missions.

In the current study program of Aerospace Engineering at Politecnico di Milano, students are introduced to the basics of control theory at the undergraduate level, where elementary concepts, from the modeling of dynamical systems in state-space form and in the transfer function domain to the design of controllers in the frequency domain, are taught in the second year course "Fundamentals of Automatic Control". At master level, the course "Aerospace Control Systems" is offered at the first year to complement the students' background with modern methods and tools for the stability and performance analysis of linear robust systems with emphasis on multi-variable Linear Time-Invariant (LTI) feedback control systems, thereby covering most of the skills needed by aerospace control engineers.

The forthcoming course "Adaptive and Autonomous aerospace Systems" (short, A&AS) to be offered to second year master's students in the first semester aims at providing the tools to face the challenges brought by the increasing demand of high-performance controllers and of unmanned vehicles capable of agile autonomous maneuvering in possibly unknown or cluttered environments. In this paper, after discussing the background in control theory of aerospace students at Politecnico di Milano, in Section 2 we present the motivation and objectives of the A&AS course. Then, in Section 3, the organization of the course is discussed and a detailed description of the covered topics is reported. The role of laboratory activities in the course will be explained in Section 4. Section 5 presents the course organization and the way in which students will be evaluated. Finally, in Section 6, we will discuss the expected learning outcomes.

Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2021.11.003

2. COURSE MOTIVATION AND OBJECTIVES

In this section we discuss the main motivations behind the introduction of the A&AS course in the master study program of Aerospace engineering at Politecnico di Milano. After an overview the students' background on automatic control, we state the objectives of the course, with focus on its contributions to the formation of future aerospace engineers.

2.1 Motivation

The most common approach to robust flight control design consists first in linearizing the aircraft model at a given flight condition and in building an uncertain model that is sufficiently representative of all the possible plant variations. Synthesis methods are then applied to find or tune a linear controller that ensures a desired level of performance while guaranteeing closed-loop stability for the uncertain plant. Extensions to achieve robustness against operating point variations are achieved by employing multi-model synthesis methods, wherein robustness and performance are guaranteed in a larger flight envelope. An alternative (less conservative) approach often employed in practice is based on the design of gain-scheduled controllers, which are robustly tuned at each operating point of interest and then blended together with a suitable scheduling algorithm. While this approach is simple, no stability guarantees are given unless the scheduling occurs at a sufficiently low rate. In this regard, Linear Parameter-Varying control (LPV control) has received a growing interest in the last two decades, providing systematic design procedures for gain-scheduled regulators in order to control dynamical systems with varying parameters which are assumed to be measurable online.

On the other hand, aerospace systems are expected to meet increasingly stringent requirements in terms of their ability to operate safely and with adequate performance in a wide range of conditions. This increasing expectation, together with the recent emergence of UAVs, is pushing towards the adoption of more advanced methods for control systems design (Lavretsky and Wise, 2013). We believe that adaptive control theory, which has reach a sound mathematical level in recent years, is the natural framework for the design of control laws for this kind of systems, which are intrinsically nonlinear, time-varying and uncertain.

As for autonomous control design, several challenges must be faced when dealing with flying vehicles (Beard et al., 2005): limited sensing capabilities and computational power pose technical difficulties for both navigation and guidance operations. For instance, when it is not possible to rely on external navigation systems (*e.g.*, the Global Navigation Satellite System), autonomous UAVs must estimate their own state through the use of onboard sensors and computing hardware. Due to weight and power limitations, aerial vehicles (and in particular small-scale UAVs) must rely on small embedded computers, lightweight laser scanners/ cameras and low-quality MEMS-based IMUs, all of which have limited ranges and field-of-view compared to their ground equivalents. As for planning, algorithms have not only to care about finding collision-free smooth paths toward a target, but they also have to deal with uncertainty in path planning and potential intermediate paths needed for re-localization. The exploration and mapping of an unknown environment can thus become a very complex decision making problem, in which a trade-off between exploration and localization of the agent must be found. At the same time, computationally efficient re-planning capabilities assume paramount importance for facing both uncertainties in the state estimate and possible moving objects. Understanding these challenges to the introduction of on-board autonomy is relevant as autonomous decision making tasks will be part, at different levels, of future air vehicles. Learning recent developments in algorithms for autonomous navigation and guidance and their interaction with autopilots for base flight control operations is also of paramount importance as autonomous control paradigms are often derived using simplified dynamical models, thereby posing stability concerns for the overall flight control scheme.

2.2 Students' background on automatic control

As mentioned in the introduction, aerospace students of Politecnico di Milano at bachelor level attend the course "Fundamentals of Automatic Control" at the second year. This course provides the students with basic notions and concepts, starting from the general definition of the control problem and the theory of dynamical systems in the time domain (state variables, system linearity, equilibria and linearization). Structural properties (stability and some elements on controllability and observability) of dynamical systems are studied as well, with emphasis on LTI systems. Laplace and Fourier transforms are then introduced: based on these tools, the description of dynamical systems in terms of transfer functions is addressed together with the analysis of the frequency response, including Bode diagrams plotting. The design of feedback control systems is addressed in terms of stability, dynamic and static performance. The design of the controller in the frequency domain is then worked out in detail, with particular reference to standard industrial controllers (PID). Root locus analysis is discussed as well. The course also provides a quick introduction to discrete time systems theory, where the main properties and the design criteria for digital control systems are discussed (digital implementation).

At the first year of master, students are offered the (optional) course "Aerospace Control Systems" (ACS), which is intended to complement the students' background with modern methods to solve aerospace control problems using state-of-the-art robust methods and tools. After reviewing basic topics in the systems theory which are only partially covered in the bachelor course (stability concepts: Lyapunov stability for equilibria of nonlinear systems; performance concepts: H_2 and H_∞ performance for linear systems), the core of ACS is on linear feedback systems. Specifically, the course syllabus covers the following topics.

- LTI Single-Input Single-Output (SISO) feedback systems (nominal design): frequency-domain loop-shaping and sensitivity shaping; time-domain state feedback and observer-based output-feedback using classical methods (eigenvalue assignment, Linear Quadratic Regulator (LQR)).
- LTI SISO feedback systems (robust design): uncertainty modeling in SISO systems; robust stability

analysis of SISO feedback systems; nominal and robust performance analysis; requirement specification; robust design: unstructured and structured mixed sensitivity synthesis.
- LTI Multiple-Input Multiple-Output (MIMO): introduction to MIMO linear systems; nominal stability and performance; robust stability and performance; robust design.

The final part of the course discusses some relevant implementation issues, notably anti-windup methods for control systems subject to saturation and the implementation of gain-scheduled controllers. Example of methods discussed during the course are then presented in the form of seminar-like lectures, where relevant applications in the context of robust control are shown.

2.3 Objectives

In view of the above discussion, the A&AS course aims at providing the students with a comprehensive understanding of the challenges related to the design of high-performance controllers for systems operating in complex environments and to give a complementary view on the design of modern feedback control systems with respect to the ACS course. Specifically, the course aims at the following goals:

- to provide a sound background on the stability analysis of nonlinear and time-varying systems;
- to introduce the concept of adaptive control, the main architectural implementations and the corresponding stability analyses;
- to present the most recent implementations of adaptive flight control systems, notably Model Reference Adaptive Control and $\mathcal{L}_1$ Adaptive Control;
- to introduce the problem of endowing aerospace systems with capabilities for autonomous operations;
- to present basic methods for autonomous guidance and navigation and their practical implementations in terms of hardware and software;
- to discuss relevant case studies in aerospace applications and implementation issues.

3. PREREQUISITES AND COURSE SYLLABUS

The course is designed to be self-contained and to be accessible to students having a bachelor in industrial engineering. A course covering the topics proposed, *e.g.*, in Fundamentals of Automatic Control (Section 2.2) should be enough to possess all the required prerequisites for an effective learning.

In greater detail, the students are required to know at least the following topics:

- fundamentals of dynamical systems; state variables, system linearity, equilibria and linearization;
- analysis of LTI systems; stability conditions, controllability and observability notions;
- design methods for MIMO LTI systems: static state feedback control design methods (eigenvalues assignment, LQR).

After an introductory lecture devoted to presenting the course motivations and to giving a general overview on adaptive control and autonomous systems with examples from flight dynamics, the syllabus of the course is split in three main parts which can be summarized as follows.

Tools for the analysis of nonlinear systems

This part is intended to provide all the necessary tools from nonlinear systems analysis which are instrumental for the comprehension of adaptive control methods. The course will present the fundamentals of nonlinear systems and phenomena with emphasis on time-varying systems. Stability definitions will be analyzed exploiting class $\mathcal{K}$ and $\mathcal{KL}$ functions. The core of this part is the Lyapunov approach to stability analysis referring specifically to nonautonomous systems. In this regard, invariance-like theorems will be presented together with stability theorems for perturbed systems by resorting to the notions of ultimate boundedness and input-to-state stability.

Adaptive control

The topic of adaptive control can be presented in different ways (Narendra and Annaswamy, 2012; Ioannou and Sun, 2012; Krstic et al., 1995): the approach that we pursue in the course follows closely the presentation of Lavretsky and Wise (2013), which is tailored to aerospace applications and contains several examples. While this reference is mostly concerned with direct Model Reference Adaptive Control (MRAC), we also devote some time discussing $\mathcal{L}_1$ adaptive control Hovakimyan and Cao (2010), which has gained attention in recent years as $\mathcal{L}_1$ controllers can be tuned to guarantee performance and robustness bounds. This part starts with an introduction to adaptive control systems, discussing the differences between robust and adaptive control architectures, direct and indirect approaches and by showing introductory examples. Then, the main results concerning the stability analysis of adaptive systems will be presented.

The Direct MRAC approach (Lavretsky and Wise, 2013) will be shown first by addressing the following topics: state feedback design for SISO and MIMO systems; adaptive augmentation of optimal baseline controllers; robust design in the presence of bounded disturbances (the dead-zone modification, the sigma-modification, the e-modification the projector operator). The presentation will be accompanied by several examples: helicopter pitch dynamics and control during hover; dynamic inversion MRAC design for scalar systems; dynamic inversion control for helicopter pitch dynamics; MRAC Control of Delta Wing Dynamics at High Angle of Attack; aircraft short-period dynamics and control. Some extensions of standard MRAC to improve transient performance will conclude this part (Closed-loop MRAC; Composite MRAC).

$\mathcal{L}_1$ adaptive control (Hovakimyan and Cao, 2010) will be introduced by showing the main difference between predictor-based MRAC formulations and $\mathcal{L}_1$ architectures. After discussing in detail the stability of the closed-loop induced by $\mathcal{L}_1$ adaptive control, a thorough performance and robustness analysis will be given with specific reference to guaranteed adaptation and robustness bounds.

The adaptive control part will be concluded by reporting some case studies, in particular, the case of adaptive augmentation for attitude and position control in multirotor

UAVs and of adaptive control for noise and vibration reduction will be shown.

Autonomous systems

The topic of autonomous systems is vast and spread across several application fields (automotive, aerospace, robotics, computer science, *etc.*). The aim of the A&AS course is to provide students with an understanding of the challenges in designing autonomous systems with specific reference to the aerospace field, give introductory notions and concepts and show basic algorithmic implementations for both navigation and guidance problems.

The first addressed topic is the problem of autonomous navigation. We will present the most common on-board sensors adopted for the autonomous navigation tasks and then discuss algorithms for Simultaneous Localization And Mapping (SLAM), Kalman filtering and graph-based methods. Then, an overview of methods for autonomous guidance in UAVs will be shown. The main topics covered by this part are path-planning and trajectory generation algorithms, real-time planning and collision avoidance methods.

As for the adaptive control part, a case study will be considered to conclude the discussion on autonomous systems. Specifically, the design of a vision-based control system for small-scale UAVs will be presented.

4. THE ROLE OF LABORATORY ACTIVITIES

As can be seen from the discussion in the previous sections, the A&AS course is fundamental in character, to allow students to be able to understand more advanced analysis and design control methods, which will become increasingly important in the near future for their careers. To compensate for the theoretical emphasis of the lectures, laboratory activities (Giurato et al., 2020; Rossiter et al., 2018) will be proposed to show both the relevance of the taught control methodologies and the challenges in actual implementations, exploiting the Flying Arena for Rotorcraft Technologies of Politecnico di Milano (FlyART, see Figure 1), a facility which has been designed to support both research and education activities in the field of multirotor UAVs, with focus on guidance, navigation and control systems. More precisely, FlyART includes an indoor fight-test facility with a $290m^3$ flight space covered by a 3D motion capture system, a few work stations for hardware integration and a classroom which can seat up to 25 students. Homemade multirotors and ANT-X drones (ANT-X, 2020) are available for flight tests in the arena. The drones (see Figure 2) are based on the open-source firmware PX4 (2020) and come with a software tool developed by ANT-X to customize the navigation and control modules of the firmware. This system allows students to develop flight control laws at high level in the MATLAB/Simulink environment, and to seamlessly deploy them in the onboard firmware, thus enabling a very short development time from model-based design environments to the actual on-board code.

Among the proposed activities, in the context of adaptive control, the design of augmented flight control systems will be presented as a laboratory lecture for both position and attitude control, leveraging the experience gathered

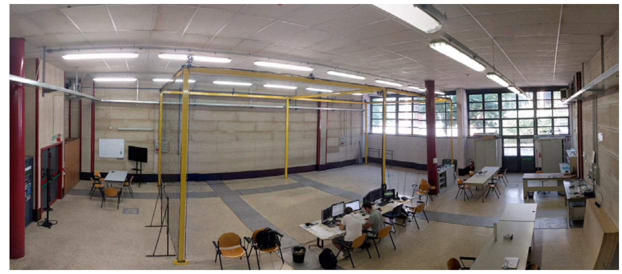

Fig. 1. The FlyART facility at Politecnico di Milano.

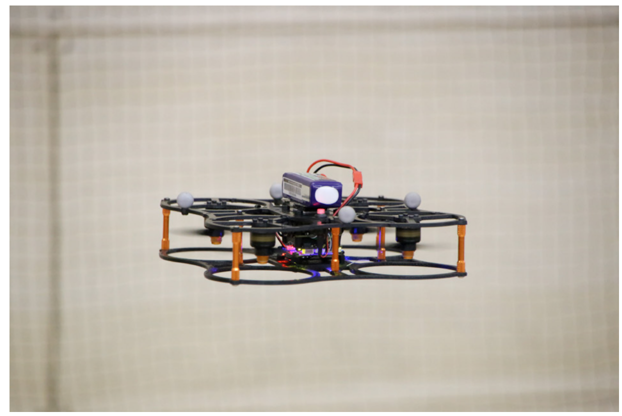

Fig. 2. The ANT-X quadrotor.

in recent research activities of the Aerospace Systems and Control Laboratory (ASCL) group of Politecnico di Milano (Bressan et al., 2020; Invernizzi et al., 2020; Invernizzi et al., 2018, 2021; Ghignoni et al., 2020). The process of developing an adaptive controller on top of an existing control architecture (adaptive augmentation) is an interesting approach to combine consolidated control strategies, which provide satisfactory results in nominal conditions, with a dynamic control law which takes care of events caused, *e.g.*, by faults, aging, environmental changes and so on. This approach, which makes the adaptive part come into play only when needed, is also interesting in view of the development of adaptive control solutions within the industrial practice. In Figure 3, one can appreciate the disturbance rejection improvement in attitude control performance obtained with MRAC and $\mathcal{L}1$ adaptive control augmentation of a robustly-tuned baseline P/PID controller, which is the most common architecture in autopilots for multirotor UAVs (see Bressan et al. (2020) for additional details). By injecting an artificial disturbance torque via software, this experiment is representative of the effect of a (partial or total) loss of throttle on one or more motors in quadrotor, as this would produce a loss of thrust, which in turns produces a disturbance torque about the body axes of the quadrotor. As for position control, the problem of wind gust rejection, which is relevant for outdoor deployment of UAVs in challenging conditions, can be addressed as well using the laboratory setup. Similar to the attitude control case shown above, the problem can be simulated by artificially injecting a suitable disturbance force to the control force computed by the controller. The results obtained in a tracking trajectory task are reported in Figure 4, where a constant disturbance is applied to the quadrotor for $5s$ while following a circular trajectory with radius $1m$ and angular rate $0.5rad/s$. The performance improvement with respect to the baseline autopilot are clearly visible, with the tracking errors reduced by more than the 50%.

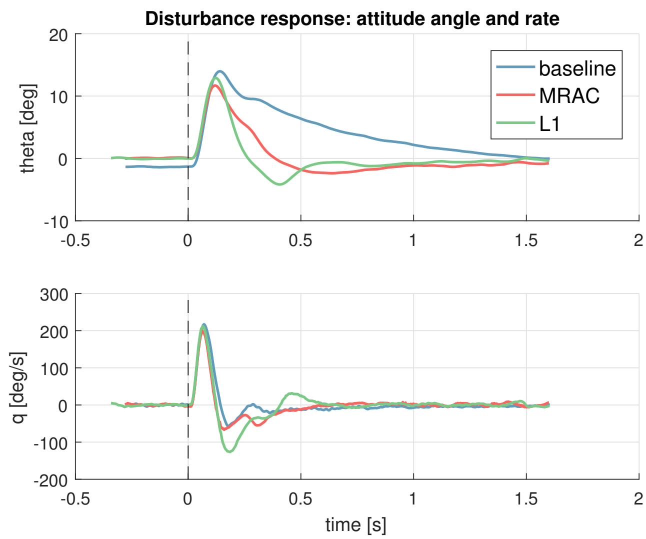

Fig. 3. Pitch angle and angular rate: response to input disturbance (comparison between baseline, MRAC and $\mathcal{L}_1$).

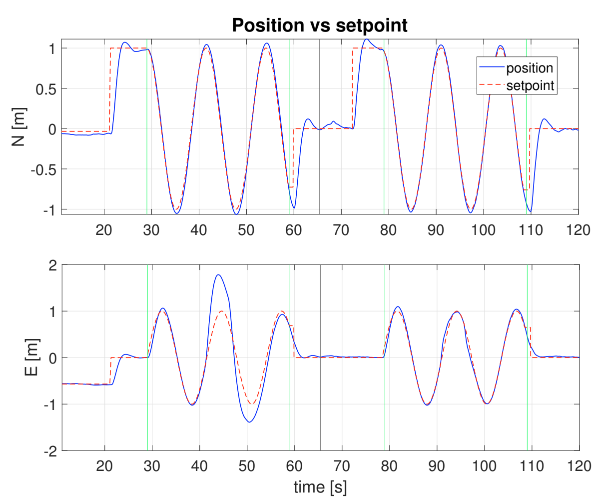

Fig. 4. North (top) and East (bottom) position of the quadrotor in the trajectory tracking experiment. Baseline controller active between 30 and 60 s, $\mathcal{L}_1$ adaptive controller active between 80 and 100 s. Constant disturbance force applied for 5 s.

For what concerns autonomous systems, the laboratory experience will address the design of a mission to explore an assigned volume in the flying arena when assuming a completely unknown environment with obstacles. The aim of the activity is to show to the students the challenges in designing an autonomous flying vehicle and how the navigation and guidance algorithms can be implemented in practice. The drone used in the mission has been developed by ANT-X (2020) and is equipped with a forward-looking (slightly pitched nose down by 15 degrees) stereo camera and a downward-looking monocular camera. For navigation, the visual odometry computed from the stereo camera images is fused with IMU information coming from the flight controller. In order to compute a globally consistent map and trajectory, a graph-based SLAM algorithm, which takes as input the stereo camera points cloud and the previously obtained odometry information, is used. The retrieved map, represented in Octomap format (Hornung et al., 2013) is then used for planning purposes (see Figure 5). At the guidance level, a real-time exploration planner is considered.

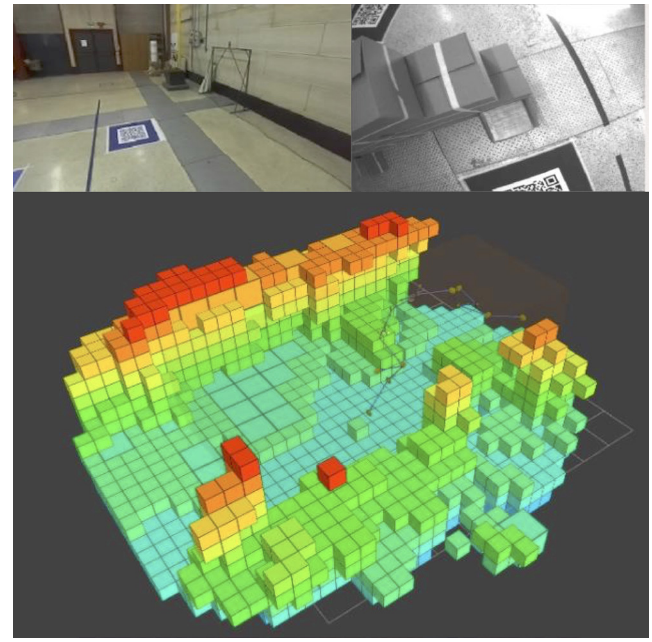

Fig. 5. Pictures collected by the onboard cameras (top), reconstructed map (bottom).

5. COURSE ORGANIZATION AND LEARNING ASSESSMENT

The A&AS course grants 6CFUs (university educational credits), where 1 CFU corresponds to 25 h of student workload (lectures plus individual study hours) and includes theoretical lectures, exercise lectures (standard and numerical) and experimental laboratory sessions, distributed across the main course topics as shown in Table 1. The total number of lecture hours is 60 and the ratio between exercise/laboratory lectures (22 h) and theoretical lectures (38 h) is about 55%. In this first implementation of the course, the adaptive control part is predominant over the autonomous one, reflecting the body of topics covered by the course (Section 3).

The learning assessment will be carried out through the assignment of a design project to small groups of students. The results must be delivered in the form of a report and an oral presentation. The latter will also include individual questions regarding the theoretical topics learned during the course. The design project will have to include: an overview of the problem; a discussion about the modeling assumptions for control design with specific reference to the uncertainties description; the choice of the control architecture; numerical results to assess the closed-loop system performance according to suitable metrics; a comparison with a classic control technique of free choice. The report will be evaluated based on methodological correctness, mathematical rigor and original contributions.

Clarity of presentation will be part of the evaluation as well. The final mark will be based on the following criteria: 50% project grade; 50% individual assessment.

Topic	Lecture	Exercise	Laboratory
Course introduction	2	0	0
Tools for the analysis of nonlinear systems	6	2	0
Adaptive control	20	10	6
Autonomous systems	10	0	4

Table 1. Syllabus for lectures (hours).

6. LEARNING OUTCOMES

The expected outcomes of the course are summarized in this section. Lectures and exercise sessions will allow students to:

- master fundamental theoretical results for nonlinear time-varying feedback control systems;
- design robust adaptive controllers for uncertain and time-varying systems with specific reference to aerospace vehicles;
- derive a suitable uncertainty model, specify desired performance and select an appropriate controller strategy;
- understand the main features and challenges of autonomous systems;
- know the necessary hardware/software components to endow UAVs with autonomous capabilities;
- learn basic navigation and guidance algorithms for autonomous flight.

7. CONCLUSIONS

In this paper we presented a new course on adaptive and autonomous systems for aerospace master students to be taught starting fall 2021 at Politecnico di Milano. We described the course motivation and objectives and presented the list of covered topics. We discussed the role of laboratory experiences that will be used to integrate the theoretical emphasis of the course, which is instrumental to provide the students with the necessary skills to master and implement advanced control methods for high performance autonomous systems operating in challenging and possibly unknown environments.

8. ACKNOWLEDGEMENTS

The Authors would like to thank Gabriele Roggi and Salvatore Meraglia for assisting in the preparation of the course material and laboratory experiences.

REFERENCES

ANT-X (2020). ANT-X software. `https://antx.it/`.

Apkarian, P., Gahinet, P., and Buhr, C. (2014). Multi-model, multi-objective tuning of fixed-structure controllers. In *2014 European Control Conference (ECC)*, 856–861.

Beard, R.W., Kingston, D., Quigley, M., Snyder, D., Christiansen, R., Johnson, W., McLain, T., and Goodrich, M. (2005). Autonomous vehicle technologies for small fixed-wing UAVs. *Journal of Aerospace Computing, Information, and Communication*, 2(1), 92–108.

Bressan, G., Russo, A., Invernizzi, D., Giurato, M., Panza, S., and Lovera, M. (2020). Adaptive augmentation of the attitude control system for a multirotor unmanned aerial vehicle. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 234(10), 1587–1596.

Dydek, Z.T., Annaswamy, A.M., and Lavretsky, E. (2013). Adaptive control of quadrotor UAVs: A design trade study with flight evaluations. *IEEE Transactions on Control Systems Technology*, 21(4), 1400–1406.

Ghignoni, P., Invernizzi, D., and Lovera, M. (2020). Fixed-dynamics antiwindup design: Application to pitch-limited position control of multirotor unmanned aerial vehicles. *IEEE Transactions on Control Systems Technology*, To appear. doi:10.1109/TCST.2020.3042073.

Giurato, M., Invernizzi, D., Panza, S., and Lovera, M. (2020). The role of laboratory activities in aerospace control education: two case studies. *IFAC-PapersOnLine*, 53(2), 17162–17167. 21th IFAC World Congress.

Hornung, A., Wurm, K.M., Bennewitz, M., Stachniss, C., and Burgard, W. (2013). OctoMap: An efficient probabilistic 3D mapping framework based on octrees. *Autonomous Robots*. `http://octomap.github.io`.

Hovakimyan, N. and Cao, C. (2010). *L1 adaptive control theory: guaranteed robustness with fast adaptation*, volume 21. Siam.

Invernizzi, D., Giurato, M., Gattazzo, P., and Lovera, M. (2018). Full pose tracking for a tilt-arm quadrotor UAV. In *Proc. IEEE Conf. Control Technology and Applications (CCTA)*, 159–164.

Invernizzi, D., Lovera, M., and Zaccarian, L. (2020). Dynamic attitude planning for trajectory tracking in thrust-vectoring UAVs. *IEEE Transactions on Automatic Control*, 65(1), 453–460.

Invernizzi, D., Giurato, M., Gattazzo, P., and Lovera, M. (2021). Comparison of control methods for trajectory tracking in fully actuated unmanned aerial vehicles. *IEEE Transactions on Control Systems Technology*, 29(3), 1147–1160.

Ioannou, P.A. and Sun, J. (2012). *Robust adaptive control*. Courier Corporation.

Johnson, E.N. and Kannan, S.K. (2005). Adaptive trajectory control for autonomous helicopters. *Journal of Guidance, Control, and Dynamics*, 28(3), 524–538.

Krstic, M., Kanellakopoulos, I., and Kokotovic, P.V. (1995). *Nonlinear and adaptive control design*. Wiley.

Lavretsky, E. and Wise, K.A. (2013). Robust and adaptive control: With aerospace applications. *Advanced textbooks in control and signal processing. London and New York: Springer*.

Mallikarjunan, S., Nesbitt, B., Kharisov, E., Xargay, E., Hovakimyan, N., Cao, C., et al. (2012). L_1 adaptive controller for attitude control of multirotors. In *AIAA Guidance, Navigation and Control Conference, Minneapolis, AIAA-2012-48312012*.

Narendra, K.S. and Annaswamy, A.M. (2012). *Stable adaptive systems*. Courier Corporation.

PX4 (2020). PX4 autopilot user guide. `https://px4.io/`.

Rossiter, J.A., Pasik-Duncan, B., Dormido, S., Vlacic, L., Jones, B., and Murray, R. (2018). A survey of good practice in control education. *European Journal of Engineering Education*, 43(6), 801–823.