

Robust Adaptive Control of Aerial Vehicles under Significant Model Uncertainty

Master's thesis in Systems, Control, and Mechatronics

Jared Alexander Tittus

DEPARTMENT OF ELECTRICAL ENGINEERING

CHALMERS UNIVERSITY OF TECHNOLOGY

Gothenburg, Sweden 2026

www.chalmers.se

MASTER'S THESIS 2026

Robust Adaptive Control of Aerial Vehicles under Significant Model Uncertainty

Jared Alexander Tittus



CHALMERS
UNIVERSITY OF TECHNOLOGY

Department of Electrical Engineering
Division of Systems and Control
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026

Robust Adaptive Control of Aerial Vehicles under Significant Model Uncertainty
Jared Alexander Tittus

© Jared Alexander Tittus, 2026.

Supervisor: Wilhelm Widlund Mellergård, SAAB Dynamics
Examiner: Balázs Kulcsár, Department of Electrical Engineering

Master's Thesis 2026
Department of Electrical Engineering
Division of Systems and Control
Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Cover: Designed control architecture for robust adaptive control. Structure is further explained in Section 3.2.

Typeset in L^AT_EX
Printed by Chalmers Reproservice
Gothenburg, Sweden 2026

Robust Adaptive Control of Aerial Vehicles under Significant Model Uncertainty
Jared Alexander Tittus
Department of Electrical Engineering
Chalmers University of Technology

Abstract

Control systems for flying vehicles must satisfy stringent demands for high performance while remaining robust against system uncertainties, such as aerodynamic variations and environmental disturbances. Conventional controllers often face a fundamental trade-off between nominal performance and robustness. This thesis investigates the robustness guarantees and performance of adaptive control methods to mitigate these challenges.

The primary focus is the design and implementation of an \mathcal{L}_1 Composite Model Reference Adaptive Controller. This architecture utilizes a cascade configuration where an outer loop employs direct MRAC to ensure tracking, while an inner loop retains an \mathcal{L}_1 adaptive structure with a fast predictor to satisfy small-gain stability conditions. The vehicle is modeled as a Linear Parameter-Varying system with decoupled dynamics for roll, pitch, and yaw.

The performance of the \mathcal{L}_1 CMRAC is systematically benchmarked against a conventional Linear Quadratic controller and a standard direct MRAC within a six-degree-of-freedom simulation environment. Evaluation is conducted across multiple reference paths designed to excite various flight dynamics and cross-coupling effects. To ensure statistical robustness, Monte Carlo simulations are utilized to quantify success rates and tracking accuracy under a broad range of uncertainty conditions. Results indicate that while the baseline LQ controller may achieve a lower root-mean-square error in nominal scenarios, the \mathcal{L}_1 CMRAC provides a higher success rate and superior robustness under significant perturbations. The analysis further highlights an inherent trade-off within the \mathcal{L}_1 framework: conservative predictor tuning enhances robustness margins but may degrade absolute tracking performance.

Keywords: Adaptive Control, \mathcal{L}_1 Adaptive Control, Robustness, Flying Vehicles, LPV Systems, Monte Carlo Simulation.

Acknowledgements

Thank you to my supervisor, Wilhelm, and my examiner, Balázs, for their support, guidance, and encouragement throughout this thesis. This collaboration has been a truly valuable experience. I also want to thank SAAB Dynamics for the opportunity and for providing the resources needed to carry out this project.

Finally, I would like to thank my friends and family for their support throughout this journey.

Jared Alexander Tittus, Gothenburg, May 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

CMRAC	Composite Model Reference Adaptive Controller
MC	Monte Carlo
MIMO	Multiple Input Multiple Output
MPC	Model Predictive Control
MRAC	Model Reference Adaptive Controller
LPV	Linear Parameter Varying
LQ/LQR	Linear Quadratic Regulator
LTI	Linear Time Invariant
PE	Persistent Exciting

Nomenclature

Below is the nomenclature of indices, sets, variables, parameters, and design parameters that have been used throughout this thesis.

Indices

i	Index of row
φ	Index for roll channel
ϑ	Index for pitch channel
ψ	Index for yaw channel
Λ	Index for input uncertainty
Ω	Index for model uncertainty
r	Index for feedforward gain
x	Index for feedback gain
l	Index for predictor model
m	Index for reference model
k	Index of discrete time step
δ	Index of control surface deflection
Δ	Index of model deviation
t	Index of time
s	Index of laplace

Sets

\mathcal{P}	Set of scheduling parameter
\mathcal{D}	Set of feasible adaptive parameter
\mathcal{S}	Set of extended feasible adaptive parameter
\mathcal{X}	Set of feasible state space defined by the imposed state constraints

Variables

$x(t)$	Measured state
$\hat{x}(t)$	Estimated state
$u(t)$	Control input
$\bar{u}(t)$	Ideal control input
$r(t)$	Reference signal
$e(t)$	Tracking error
$\tilde{x}(t)$	Prediction error
$V(x, t)$	Lyapunov candidate function
$\Lambda(t)$	True input uncertainty
$\hat{\Lambda}(t)$	Estimated input uncertainty
$\tilde{\Lambda}(t)$	Estimation Error of input uncertainty
$\Lambda(t)$	Input uncertainty
$\Delta(x, t)$	True model uncertainty
$\Theta(t)$	True uncertainty weights
$\hat{\Theta}(t)$	Estimated uncertainty weights
$\tilde{\Theta}(t)$	Estimation error of uncertainty weights
$\Omega(t)$	True scaled uncertainty weights
$\hat{\Omega}(t)$	Estimated scaled uncertainty weights
$\tilde{\Omega}(t)$	Estimation error scaled uncertainty weights
$\hat{K}(t)$	Estimated control gain
$\tilde{K}_r(t)$	Estimation error of control gain
$\rho(t)$	Varying parameter
$A(\rho, t)$	Linear parameter varying plant matrix
$B(\rho, t)$	Linear parameter varying input matrix
$C(\rho, t)$	Linear parameter varying output matrix
$D(\rho, t)$	Linear parameter varying feed-through matrix
$P(\rho, t)$	Linear parameter varying Lyapunov solution
$G(s)$	Transfer function

Parameters

A	Linear time invariant plant matrix
B	Linear time invariant input matrix
C	Linear time invariant output matrix
D	Linear time invariant feed-through matrix
K	Linear time invariant control gain
L	Linear time invariant predictor gain
P	Linear time invariant Lyapunov solution

Design Parameters

Γ	Learning rate
Q	State penalty
η	Bound scaling factor
ϵ	Projection tolerance
$K(\rho, t)$	Scheduled reference control gain
$L(\rho, t)$	Scheduled predictor gain
$\omega_c(\rho, t)$	Scheduled cut-off frequency
$\Phi(x, t)$	Lipschitz-continuous basis functions



Contents

List of Acronyms	ix
Nomenclature	xi
List of Figures	xix
List of Tables	xxiii
1 Introduction	1
1.1 Background	1
1.2 Purpose	2
1.3 Scope	2
1.4 Objective / Research Questions	3
1.5 Ethics & Sustainability	3
1.6 Delimitations	3
2 Theory	5
2.1 Lyapunov Stability	5
2.1.1 Lyapunov's Direct Method	5
2.1.2 Quadratic Lyapunov Functions for Linear Systems	5
2.2 Model Reference Adaptive Control	6
2.2.1 Uncertainty Model	6
2.2.2 Direct MRAC	7
2.2.3 Indirect MRAC	8
2.2.4 Persistent Exciting	10
2.3 Robust Control Theory	11
2.3.1 Norms	11
2.3.2 Small-Gain Theorem	13
2.4 \mathcal{L}_1 Adaptive Control	13
2.5 Projection Operator	15
2.5.1 Convex Set and Smooth Convex Function	15
2.5.2 Projection Operator Definition	16
2.5.3 Key Properties	16
2.6 Linear Quadratic Regulator	17
2.6.1 Linear Parameter-Varying Systems	18
2.6.2 Gain Scheduling	19
2.7 Parameter Summary	19

3	Methods	21
3.1	System Description	21
3.1.1	Vehicle Structure	21
3.1.2	Simulation Model	22
3.1.3	Baseline Controller	24
3.1.4	Guidance Module	25
3.2	Design of Robust Adaptive Control	26
3.2.1	Inner Loop	26
3.2.2	Outer Loop	27
3.2.3	Extended Projection Operator	28
3.3	Design of Reference Model	28
3.4	Design of Prediction Model	29
3.5	Implemented \mathcal{L}_1 CMRAC Architecture	31
3.5.1	Tuning Parameters	32
3.6	Evaluation	34
3.6.1	Paths	34
3.6.2	Predictor Evaluation	36
3.6.3	Monte Carlo	37
3.6.4	Sensitivity Analysis	39
4	Results	41
4.1	Scaling Factor Tuning	41
4.2	\mathcal{L}_1 Norm	42
4.3	Controller Performance	44
4.3.1	Path 1	44
4.3.2	Path 2	55
4.3.3	Path 3	66
4.4	Controller Performance in the Presence of External Uncertainty	77
5	Discussion	79
5.1	Importance of Scaling Factor η for Input Uncertainty Bound Λ_{\max}	79
5.2	Validation of \mathcal{L}_1 Norms	79
5.3	Overall Performance Comparison	80
5.3.1	Path Dependent Performance	81
5.4	External Disturbance Attenuation	82
5.5	Accuracy–Robustness Trade-off	83
5.6	Practical Comparison	85
6	Conclusion	87
6.1	Question 1: Robustness Guarantees	87
6.2	Question 2: Balance Between Robustness and Performance	88
6.3	Future Work	88
	Bibliography	91

A	Appendix 1	I
A.1	Evaluation of the \mathcal{L}_1 CMRAC Architecture Using a Constant-Velocity Model	I

List of Figures

2.1	Block diagram of the direct MRAC controller. The control input u is computed using the adaptive feedback and feedforward gains, \hat{K}_x and \hat{K}_r , and the estimated uncertainty $\hat{\Theta}$. The reference model (2.7) provides the desired state trajectory x_m , which, together with the plant (2.6) output x , generates the tracking error e . The adaptive update laws adjust \hat{K}_x , \hat{K}_r , and $\hat{\Theta}$ based on the prediction error to ensure boundedness of the error and convergence of the system states to the reference model.	9
2.2	Block diagram of the indirect MRAC controller. The control input u is generated using the nominal feedforward gain K_r and the estimated uncertainties $\hat{\Lambda}^{-1}$ and $\hat{\Omega}$, which are computed by the adaptive update laws. The predictor model (2.22) estimates the system states \hat{x} , which together with the plant (2.21) output x forms the prediction error \tilde{x} . This error is used to update the adaptive parameters $\hat{\Lambda}$ and $\hat{\Omega}$	10
2.3	Generalized robust control interconnection. The generalized plant is partitioned to expose the uncertainty channel (u_Δ, y_Δ) and the performance channel (w, z). The controller K generates the control input u from the measured output v . The block Δ represents structured or unstructured uncertainty in feedback with the nominal closed-loop system. Robust stability and performance are assessed with respect to this interconnection.	11
2.4	Block diagram of the \mathcal{L}_1 adaptive control architecture. The control input u consists of the nominal feedforward term $K_r^\top r$ and the adaptive component filtered through the strictly proper low-pass filter $C(s)$, i.e., $-C(s)\hat{\Lambda}^{-1}\hat{\Omega}^\top\Phi(x)$. The state predictor (2.22) generates the estimated states \hat{x} based on the measured plant state x from (2.21) and the control input. The prediction error $\tilde{x} = \hat{x} - x$ is used in the adaptive laws to update the uncertainty estimates $\hat{\Omega}$ and $\hat{\Lambda}$	14
2.5	Block diagram of the LQ state-feedback controller with reference feedforward gain. The control input is given by $u = K_r r - K_x x$, where K_x is the optimal state-feedback gain obtained by minimizing a quadratic cost function, and K_r is the reference feedforward gain ensuring steady-state tracking of the reference signal r . The states x from the plant (2.57) is fed back through the gain K_x to regulate the system states and achieve optimal performance in the sense of the chosen quadratic cost functional.	18

3.1	Illustration of Aerial Vehicle. Figure is obtained from [1].	21
3.2	Polytope of operating-points for model linearization.	22
3.3	Polytope of operating-points for gain scheduling.	24
3.4	Example of guidance module commanding a vehicle	26
3.5	2D illustration of row-wise projection of the parameter update onto the ℓ_1 -ball in the extended projection operator. The red line represents the correct update step	29
3.6	Block diagram of the \mathcal{L}_1 composite MRAC architecture. The plant dynamics in (2.6) are controlled to follow the reference model in (2.7). The control input (3.9) includes a strictly proper low-pass filter $C(s)$. The predictor (3.10) generates the estimate \hat{x} , forming the prediction error \tilde{x} , while the tracking error e is computed relative to the reference model. Both errors drive the composite adaptive laws, and the predictor dynamics are selected to satisfy the small-gain condition in (2.48), ensuring bounded-input bounded-output stability.	32
3.7	Baseline evaluation for Path 1.	35
3.8	Baseline evaluation for Path 2.	35
3.9	Baseline evaluation for Path 3.	36
3.10	\mathcal{L}_1 norms produced with predictor tuned for nominal Ω_{\max}	37
3.11	Monte Carlo traces of the integral states for Path 1 obtained from 500 simulations using the LQ controller.	38
3.12	Monte Carlo traces of the integral states for Path 2 obtained from 500 simulations using the LQ controller.	39
3.13	Monte Carlo traces of the integral states for Path 3 obtained from 500 simulations using the LQ controller.	39
3.14	Applied disturbance perpendicular to velocity vector of vehicle.	40
4.1	Adaptive parameters for different scaling factors η on Path 1.	41
4.2	\mathcal{L}_1 norms produced with predictor tuned for $1.2 \times \Omega_{\max}$	42
4.3	\mathcal{L}_1 norms produced with predictor tuned for $0.7 \times \Omega_{\max}$	43
4.4	Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the LQ controller.	44
4.5	Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the LQ controller.	45
4.6	Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the MRAC controller.	46
4.7	Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the MRAC controller.	47
4.8	Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.	48
4.9	Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.	49
4.10	Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.	50

4.11	Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.	51
4.12	Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.	52
4.13	Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.	53
4.14	Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the LQ controller.	55
4.15	Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the LQ controller.	56
4.16	Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the MRAC controller.	57
4.17	Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the MRAC controller.	58
4.18	Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.	59
4.19	Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.	60
4.20	Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.	61
4.21	Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.	62
4.22	Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.	63
4.23	Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.	64
4.24	Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the LQ controller.	66
4.25	Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the LQ controller.	67
4.26	Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the MRAC controller.	68
4.27	Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the MRAC controller.	69
4.28	Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.	70
4.29	Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.	71
4.30	Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.	72

4.31	Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.	73
4.32	Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.	74
4.33	Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.	75
4.34	Adaptive parameter evolution under external disturbance for Path 3. The disturbance is active between $t = 22$ s and $t = 25$ s.	77
4.35	Closed-loop response under external disturbance for Path 3. The figure shows the roll angle and acceleration tracking for the LQ, MRAC, and \mathcal{L}_1 CMRAC controllers, together with the corresponding nominal trajectories. The disturbance is active between $t = 22$ s and $t = 25$ s.	78
A.1	Monte Carlo trajectory realizations for the \mathcal{L}_1 CMRAC and LQ controllers applied to the 1D constant-velocity model.	II

List of Tables

2.1	Parameters and their dependencies	20
3.1	Measured States	23
3.2	Tuning parameters and their value for roll channel	33
3.3	Tuning parameters and their value for pitch and yaw channel	33
3.4	A priori success rate estimated from initial Monte Carlo simulations.	38
4.1	Performance results of each controller for roll channel in Path 1. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.	54
4.2	Performance results of each controller for pitch channel in Path 1. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.	54
4.3	Performance results of each controller for yaw channel in Path 1. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.	54
4.4	Performance results of each controller for roll channel in Path 2. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.	65
4.5	Performance results of each controller for pitch channel in Path 2. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.	65
4.6	Performance results of each controller for yaw channel in Path 2. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.	65
4.7	Performance results of each controller for roll channel in Path 3. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.	76
4.8	Performance results of each controller for pitch channel in Path 3. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.	76
4.9	Performance results of each controller for yaw channel in Path 3. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.	76

5.1	Summarized success rates with controllers ranked from highest to lowest based on average success rate. ST, NT, and LT denote the predictor tuning for small, nominal, and large uncertainty bound, respectively.	80
A.1	Performance results of each controller for position. RMSE and MAE are averaged over all MC simulations.	II
A.2	Performance results of each controller for velocity. RMSE and MAE are averaged over all MC simulations.	II

1

Introduction

This chapter provides an overview of the thesis and establishes the context for the work. It begins by introducing the background and motivation for studying robust and adaptive control of flying vehicles under uncertainty. The purpose and scope of the thesis are then defined, followed by the formulation of the research objectives and questions that guide the study. Finally, ethical considerations and project delimitations are discussed to clarify the boundaries and context of the work.

1.1 Background

The control systems of flying vehicles are required to satisfy stringent demands for both high performance and robustness, particularly in the presence of system uncertainties and operational variations. Such uncertainties may arise from manufacturing tolerances, differences in material properties, aerodynamic parameter variations, or environmental disturbances such as wind gusts and turbulence. In conventional control design, achieving both high performance and robustness often involves a fundamental trade-off: controllers optimized for nominal performance can be highly sensitive to deviations in system parameters, whereas robust controllers provide stability across a range of uncertainties but may reduce overall system performance. In conventional control design, a common mitigation strategy is to impose strict requirements on manufacturing precision, as small variations between vehicles can significantly affect the ability to maintain both performance and robustness. Consequently, ensuring uniformity across units often increases production costs and can hinder the economic feasibility of mass production.

Adaptive control methods offer a promising alternative by allowing the control law to adjust online in response to variations in system dynamics. Unlike fixed-gain controllers, adaptive controllers can modify their behavior based on observed performance, enabling more consistent operation across vehicles with varying parameters. This capability is particularly advantageous when there are large uncertainties in system models, where conventional fixed controllers would either be too conservative or fail to meet performance specifications. However, adaptive methods also introduce challenges. Because the control law evolves during operation, traditional robustness guarantees are typically weaker compared to non-adaptive designs. In safety-critical applications such as aviation, where stability and performance are paramount, this limitation presents a significant concern and necessitates careful analysis to ensure that adaptive controllers remain reliable under all realistically possible conditions.

1.2 Purpose

The purpose of this thesis is to investigate the robustness guarantees of adaptive control methods when applied to systems with uncertain parameters. Specifically, certain system parameters are assumed to be stochastic, with known probability distributions defined within bounded intervals. The primary focus is on identifying and evaluating different control methodologies that can maintain robustness across the full range of possible system realizations induced by these uncertainties. By evaluating adaptive controllers in this context, the study aims to understand the extent to which robustness can be preserved while still leveraging the performance advantages of adaptation.

The work will be conducted within a six-degree-of-freedom simulation environment containing a flying vehicle subject to significant model uncertainties. As the simulation environment has been used before, this work can be viewed as a continuation of [1], which incorporated and evaluated adaptive control laws for an aerial vehicle within the same simulation environment used. For this thesis, an adaptive control law will be designed and implemented within this environment, and its robustness and performance will be systematically analyzed. The evaluation will include comparisons with conventional control strategies, including classical linear-quadratic controllers as well as other adaptive methods, to benchmark performance and quantify robustness under nominal and perturbed operating conditions. Through this analysis, the thesis seeks to provide insight into the design of robust adaptive controllers suitable for safety-critical, uncertain systems, highlighting the potential advantages and limitations of adaptive methods relative to traditional control approaches.

1.3 Scope

The scope of this thesis is confined to the design and analysis of adaptive and robust control methods for a flying vehicle in a six-degree-of-freedom simulation environment. The evaluated controller is based on the \mathcal{L}_1 adaptive controller proposed by [2].

Performance and robustness will be assessed primarily by applying reference accelerations or commands to the vehicle and examining whether the control system satisfies predefined stability and tracking requirements across the defined uncertainty set. Both nominal and perturbed scenarios will be considered to evaluate the consistency and reliability of each controller's performance.

The thesis is restricted entirely to simulation studies. No hardware testing, real-time implementation, or physical validation is included. The scope is further limited to exclude practical considerations such as sensor noise modeling, actuator saturation, and aerodynamic refinements, unless they directly influence the evaluation of robustness under parametric uncertainty. All stochastic parameters are assumed to follow known probability distributions within predefined bounded intervals.

1.4 Objective / Research Questions

The primary objective of this thesis is to investigate the robustness guarantees of adaptive control laws applied to flying vehicles with uncertain system parameters. Specifically, the research seeks to answer the following questions:

1. **Robustness Guarantees:** What robustness guarantees can be provided for the designed adaptive controller? This will be compared to the robustness of a conventional LQ controller designed for the nominal system, evaluated using statistical methods.
2. **Balance Between Robustness and Performance:** Is it necessary, even with adaptive control, to balance robustness against performance? This investigation will assess whether the conventional trade-off between robustness and performance also applies to adaptive control methods, and if so, to what extent.

1.5 Ethics & Sustainability

This project is conducted in collaboration with SAAB Dynamics, a company operating in the defense sector. As such, ethical considerations regarding the potential impact and application of the research must be addressed. Research in this area may involve dual-use technologies, which can have both civilian and defense-related applications.

The work presented in this thesis is strictly limited to a simulation-based case study. No results will be deployed, tested, or directly implemented in operational or safety-critical systems. The primary aim is to investigate theoretical aspects of adaptive and robust control under model uncertainty, contributing to academic knowledge rather than the development of weapon systems or autonomous decision-making capabilities.

SAAB Dynamics is a technologically innovative company that encourages creative and advanced engineering solutions. The company also writes regarding sustainability: "*Peace and security are fundamental for human rights and freedom. We contribute to sustainable development by delivering capabilities that allow our customers to protect people and society.*" [3]. The collaboration provides an opportunity to study complex control problems in an industrial context while maintaining clear academic and ethical boundaries.

From a societal perspective, maintaining technological competence in national defense is an important aspect of ensuring security and stability. In the current geopolitical context, research that improves the robustness and reliability of control systems can be viewed as contributing to defensive preparedness rather than offensive applications.

1.6 Delimitations

Due to the collaboration with the defense sector, certain information regarding the simulation environment will be omitted from this report. These include the vehicle

1. Introduction

model, and information about the baseline controllers.

2

Theory

This chapter presents the theoretical frameworks used to derive a robust adaptive controller.

Unless otherwise stated, all signals in this report are functions of time. To avoid notational clutter, the explicit time argument (t) is omitted throughout the thesis except equations which include integrals, and each signal should be understood as time-varying. Equivalent representations of the same signal are used interchangeably when no ambiguity arises.

Table 2.1 summarizes all parameters used in the following chapter to further clarify their dependencies.

2.1 Lyapunov Stability

Stability analysis is a fundamental aspect of control design, ensuring that the system trajectories remain bounded and converge to desired states. In adaptive control Lyapunov-based methods are the standard tool for establishing stability and deriving adaptation laws.

2.1.1 Lyapunov's Direct Method

Lyapunov's direct method provides a constructive way to assess stability without explicitly solving the system dynamics. It involves selecting a scalar function $V(x)$, called a Lyapunov candidate function, which satisfies the following conditions in a neighborhood of the equilibrium $x = 0$:

$$V(x) > 0, \quad \forall x \neq 0, \quad V(0) = 0, \quad (2.1)$$

$$\dot{V}(x) = \frac{\partial V}{\partial x} f(x, t) \leq 0, \quad \forall x. \quad (2.2)$$

If such a $V(x)$ exists, the equilibrium at $x = 0$ is stable. If $\dot{V}(x) < 0$ for all $x \neq 0$, the equilibrium is asymptotically stable [4].

2.1.2 Quadratic Lyapunov Functions for Linear Systems

For a linear time-invariant system

$$\dot{x} = Ax, \quad (2.3)$$

given that A is Hurwitz [5], a common choice is a quadratic Lyapunov function

$$V(x) = x^\top P x, \quad (2.4)$$

where $P = P^\top > 0$ satisfies the Lyapunov equation

$$A^\top P + P A = -Q, \quad (2.5)$$

with $Q = Q^\top > 0$ chosen by the designer. The existence of such a P guarantees that the system is exponentially stable [4].

2.2 Model Reference Adaptive Control

Model Reference Adaptive Control (MRAC) is an adaptive control strategy in which the controller parameters are adjusted online such that the system state follows a desired reference model despite parametric uncertainties.

Consider the uncertain system

$$\dot{x} = A x + B \Lambda (u + \Delta(x)), \quad (2.6)$$

where $x \in \mathbb{R}^n$ is the state vector and $u \in \mathbb{R}^m$ is the control input. The matrices $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$ represent the known nominal plant dynamics where (A, B) is controllable [6]. The matrix $\Lambda \in \mathbb{R}^{m \times m}$ models unknown input gain uncertainty and is assumed non-singular. The matrix $\Delta(x) \in \mathbb{R}^m$ contains unknown constant parameters associated with matched uncertainties [7]. In this thesis, the uncertainties Λ and $\Delta(x)$ are assumed to be time dependent. However, to maintain concise notation, the explicit time argument (t) is omitted.

The desired closed-loop behavior is specified by the reference model

$$\dot{x}_m = A_m x_m + B_m r, \quad (2.7)$$

where $x_m \in \mathbb{R}^n$ is the reference state and $r \in \mathbb{R}^m$ is a bounded reference input. The matrix $A_m \in \mathbb{R}^{n \times n}$ is chosen Hurwitz to ensure exponential stability of the reference dynamics, while $B_m \in \mathbb{R}^{n \times m}$ determines the steady-state gain and tracking characteristics.

2.2.1 Uncertainty Model

The model uncertainty $\Delta(x)$ can be modeled as the linear combination between N known, locally Lipschitz-continuous basis functions $\varphi_i(x)$ with unknown coefficients [8]. A Lipschitz-continuous function satisfies the condition

$$|\varphi_i(x_1) - \varphi_i(x_2)| \leq L_i |x_1 - x_2|, \quad i \in \{1, \dots, N\} \quad (2.8)$$

for some constant $L_i > 0$ and for all x_1, x_2 in the domain. This condition restricts the rate of variation of the function and ensures that small changes in the state produce proportionally bounded changes in the basis functions [4]. Thus, $\Delta(x)$ can be rewritten to the form

$$\Delta(x) = \Theta^\top \Phi(x) \quad (2.9)$$

where $\Phi(x) \in \mathbb{R}^N$ and $\Theta \in \mathbb{R}^{N \times m}$.

A linear system affected by parametric uncertainty can generally be written as

$$\dot{x} = (A + A_\Delta)x + (B + B_\Delta)u, \quad (2.10)$$

where A_Δ and B_Δ denote unknown deviations from the nominal model. If the uncertainty is matched, i.e., enters through the input channel, (2.10) can be expressed in the form of (2.6). Assuming that B has full column rank, the uncertainty can be parameterized using the Moore–Penrose pseudoinverse B^\dagger , [9], as

$$\Theta = (B + B_\Delta)^\dagger A_\Delta, \quad \Lambda = I + B^\dagger B_\Delta. \quad (2.11)$$

2.2.2 Direct MRAC

The MRAC controller can be divided into two categories: direct and indirect MRAC. In direct MRAC, the control law is designed directly to match the system to a reference model without explicitly estimating the plant parameters.

Assume the reference model satisfies

$$A_m x_m + B_m r = Ax + B\Lambda(\bar{u} + \Theta^\top \Phi(x)), \quad (2.12)$$

where

$$\bar{u} = K_r^\top r - K_x^\top x - \Theta^\top \Phi(x) \quad (2.13)$$

is the ideal control law. Here, K_r , K_x , and Θ denote the true feedforward gain, feedback gain, and system uncertainty, respectively.

Using (2.12), the reference model matrices can be related to the true plant parameters as

$$A_m = A - B\Lambda K_x, \quad B_m = B\Lambda K_r. \quad (2.14)$$

The applied adaptive control law is then

$$u = \hat{K}_r r - \hat{K}_x x - \hat{\Theta}^\top \Phi(x), \quad (2.15)$$

where \hat{K}_r , \hat{K}_x , and $\hat{\Theta}$ are the adaptive gains updated online based on the tracking error

$$\begin{aligned} e &= x - x_m, \\ \dot{e} &= Ax + B\Lambda(u + \Theta^\top \Phi(x)) - A_m x_m - B_m r. \end{aligned} \quad (2.16)$$

Substituting (2.14) and (2.15) into (2.16), the tracking error dynamics can be expressed as

$$\dot{e} = A_m e + B\Lambda(\tilde{K}_r^\top r - \tilde{K}_x^\top x - \tilde{\Theta}^\top \Phi(x)), \quad (2.17)$$

where $\tilde{K}_r = \hat{K}_r - K_r$, $\tilde{K}_x = \hat{K}_x - K_x$, and $\tilde{\Theta} = \hat{\Theta} - \Theta$ denote the parameter estimation errors.

To guarantee stability for the tracking error, the following Lyapunov candidate is proposed [8]:

$$V(e, \tilde{K}_r, \tilde{K}_x, \tilde{\Theta}) = e^\top P e + \text{tr} \left([\tilde{K}_r^\top \Gamma_r^{-1} \tilde{K}_r + \tilde{K}_x^\top \Gamma_x^{-1} \tilde{K}_x + \tilde{\Theta}^\top \Gamma_\Theta^{-1} \tilde{\Theta}] \Lambda \right), \quad (2.18)$$

where $P = P^\top > 0$ solves the Lyapunov equation for the reference model, Γ_r , Γ_x , and Γ_Θ are positive definite learning rate matrices, and $\text{tr}(\cdot)$ denotes the trace operator (sum of diagonal elements) [4].

The derivative of the Lyapunov function is

$$\begin{aligned} \dot{V} &= e^\top P \dot{e} + \dot{e}^\top P e + 2 \text{tr} \left(\left[\tilde{K}_r^\top \Gamma_r^{-1} \dot{\tilde{K}}_r + \tilde{K}_x^\top \Gamma_x^{-1} \dot{\tilde{K}}_x + \tilde{\Theta}^\top \Gamma_\Theta^{-1} \dot{\tilde{\Theta}} \right] \Lambda \right) \\ &= -e^\top Q e + 2 \text{tr} \left(\tilde{K}_r^\top \left[\Gamma_r^{-1} \dot{\tilde{K}}_r + r e^\top P B \right] \Lambda \right) + 2 \text{tr} \left(\tilde{K}_x^\top \left[\Gamma_x^{-1} \dot{\tilde{K}}_x - x e^\top P B \right] \Lambda \right) \\ &\quad + 2 \text{tr} \left(\tilde{\Theta}^\top \left[\Gamma_\Theta^{-1} \dot{\tilde{\Theta}} - \Phi(x) e^\top P B \right] \Lambda \right), \end{aligned} \tag{2.19}$$

where $A_m^\top P + P A_m = -Q$.

Choosing the adaptive update laws as

$$\begin{aligned} \dot{\tilde{K}}_r &= -\Gamma_r r e^\top P B, \\ \dot{\tilde{K}}_x &= \Gamma_x x e^\top P B, \\ \dot{\tilde{\Theta}} &= -\Gamma_\Theta \Phi(x) e^\top P B, \end{aligned} \tag{2.20}$$

ensures that $\dot{V} \leq 0$, which guarantees that the tracking error and parameter estimation errors remain bounded, and the closed-loop system is stable. The control synthesis is summarized by Figure 2.1

2.2.3 Indirect MRAC

In contrast to direct MRAC, where the controller parameters are adapted directly, indirect MRAC introduces an intermediate parameter estimation step. The unknown plant parameters are first identified online using a state predictor, and the control law is subsequently constructed using the parameter estimates.

Consider the uncertain plant model

$$\dot{x} = A_m x + B \Lambda \left(u + \Theta^\top \Phi(x) \right), \tag{2.21}$$

where A_m is a known Hurwitz reference matrix defined by (2.14).

The predictor model is defined as

$$\dot{\hat{x}} = A_m \hat{x} + B \left(\hat{\Lambda} u + \hat{\Omega}^\top \Phi(x) \right), \quad \Omega = \Lambda \Theta, \tag{2.22}$$

where $\hat{\Lambda}$ and $\hat{\Omega}$ denote the parameter estimates.

The control law is constructed using the parameter estimates as

$$u = K_r r - \hat{\Lambda}^{-1} \hat{\Theta}^\top \Phi(x), \tag{2.23}$$

where K_r is a nominal feedforward gain ensuring reference tracking, and the nominal feedback dynamics are embedded in A_m to guarantee desired stability properties [2]. In indirect MRAC, the adaptive mechanism estimates the plant uncertainty rather than directly updating the control gains.

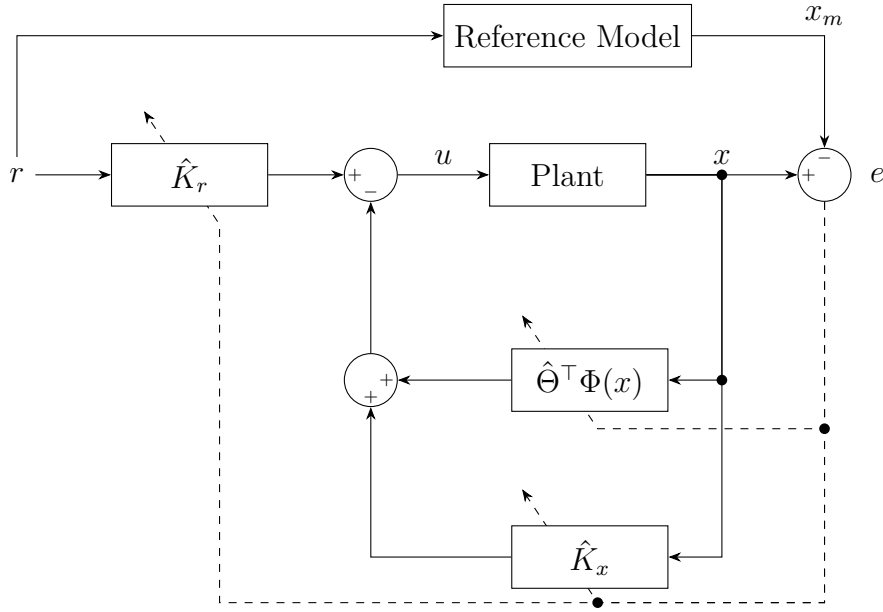


Figure 2.1: Block diagram of the direct MRAC controller. The control input u is computed using the adaptive feedback and feedforward gains, \hat{K}_x and \hat{K}_r , and the estimated uncertainty $\hat{\Theta}$. The reference model (2.7) provides the desired state trajectory x_m , which, together with the plant (2.6) output x , generates the tracking error e . The adaptive update laws adjust \hat{K}_x , \hat{K}_r , and $\hat{\Theta}$ based on the prediction error to ensure boundedness of the error and convergence of the system states to the reference model.

Rather than updating the parameter estimates via the tracking error (2.16), the adaptive laws are derived via the prediction error

$$\tilde{x} = \hat{x} - x. \quad (2.24)$$

where the error dynamics become

$$\begin{aligned} \dot{\tilde{x}} &= A_m \hat{x} + B (\hat{\Lambda} u + \hat{\Omega}^T \Phi(x)) - A_m x - B \Lambda (u + \Theta^T \Phi(x)) \\ &= A_m \tilde{x} + B (\tilde{\Lambda} u + \tilde{\Omega}^T \Phi(x)), \end{aligned} \quad (2.25)$$

Here, the parameter estimation errors are defined as

$$\tilde{\Lambda} = \hat{\Lambda} - \Lambda, \quad \tilde{\Omega} = \hat{\Omega} - \Omega. \quad (2.26)$$

To establish boundedness of the estimation errors, consider the Lyapunov candidate function

$$V(\tilde{x}, \tilde{\Lambda}, \tilde{\Omega}) = \tilde{x}^T P \tilde{x} + \text{tr}(\tilde{\Lambda}^T \Gamma_\Lambda^{-1} \tilde{\Lambda} + \tilde{\Omega}^T \Gamma_\Omega^{-1} \tilde{\Omega}), \quad (2.27)$$

where $P = P^T > 0$ satisfies the Lyapunov equation (2.5) for $A = A_m$ and $\Gamma_\Lambda, \Gamma_\Omega > 0$ are adaptation gain matrices.

Differentiating V along the trajectories of the error dynamics yields

$$\begin{aligned} \dot{V} &= \tilde{x}^T (A_m^T P + P A_m) \tilde{x} + 2 \tilde{x}^T P B (\tilde{\Lambda} u + \tilde{\Omega}^T \Phi(x)) \\ &\quad + 2 \text{tr}(\tilde{\Lambda}^T \Gamma_\Lambda^{-1} \dot{\tilde{\Lambda}} + \tilde{\Omega}^T \Gamma_\Omega^{-1} \dot{\tilde{\Omega}}). \end{aligned} \quad (2.28)$$

where using (2.5) this becomes

$$\begin{aligned} \dot{V} = & -\tilde{x}^\top Q \tilde{x} + 2 \operatorname{tr} \left(\tilde{\Lambda}^\top \left[\Gamma_\Lambda^{-1} \dot{\hat{\Lambda}} + u \tilde{x}^\top P B \right] \right) \\ & + 2 \operatorname{tr} \left(\tilde{\Omega}^\top \left[\Gamma_\Omega^{-1} \dot{\hat{\Omega}} + \Phi(x) \tilde{x}^\top P B \right] \right). \end{aligned} \quad (2.29)$$

Choosing the adaptive laws as

$$\begin{aligned} \dot{\hat{\Omega}} &= -\Gamma_\Omega \Phi(x) \tilde{x}^\top P B, \\ \dot{\hat{\Lambda}} &= -\Gamma_\Lambda u \tilde{x}^\top P B, \end{aligned} \quad (2.30)$$

cancel the cross terms and yields

$$\dot{V} = -\tilde{x}^\top Q \tilde{x} \leq 0. \quad (2.31)$$

satisfying (2.2). Hence, \tilde{x} , $\tilde{\Lambda}$, and $\tilde{\Omega}$ are bounded, and the prediction error \tilde{x} converges to zero asymptotically. The control architecture is summarized in Figure 2.2 where the closed-loop plant is defined by (2.21), i.e. closed-loop stable.

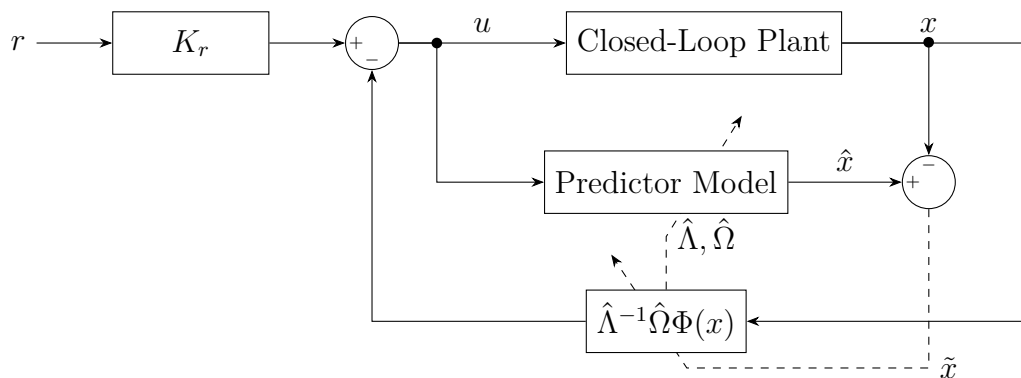


Figure 2.2: Block diagram of the indirect MRAC controller. The control input u is generated using the nominal feedforward gain K_r and the estimated uncertainties $\hat{\Lambda}^{-1}$ and $\hat{\Omega}$, which are computed by the adaptive update laws. The predictor model (2.22) estimates the system states \hat{x} , which together with the plant (2.21) output x forms the prediction error \tilde{x} . This error is used to update the adaptive parameters $\hat{\Lambda}$ and $\hat{\Omega}$.

2.2.4 Persistent Exciting

While the Lyapunov analysis guarantees boundedness of the parameter estimation errors, convergence of the parameter estimates to their true values requires additional excitation conditions. A regressor signal $\Phi(x)$ is said to be persistently exciting (PE) [10], if there exist constants $\gamma, T > 0$ such that

$$\int_t^{t+T} \Phi(\tau) \Phi^\top(\tau) d\tau \geq \gamma I, \quad \forall t \geq 0. \quad (2.32)$$

Under persistent excitation, the parameter estimation errors $\tilde{\Lambda}$ and $\tilde{\Omega}$ converge exponentially to zero. In the absence of persistent excitation, parameter estimates remain bounded but may not converge to their true values.

2.3 Robust Control Theory

Robust control theory concerns the analysis and synthesis of feedback systems that guarantee internal stability and prescribed performance in the presence of structured or unstructured uncertainty [4]. In contrast to nominal control design, which assumes an exact mathematical description of the plant, robust control explicitly models uncertainty and derives conditions under which closed-loop properties are preserved for all admissible perturbations.




Figure 2.3: Generalized robust control interconnection. The generalized plant is partitioned to expose the uncertainty channel (u_Δ, y_Δ) and the performance channel (w, z) . The controller K generates the control input u from the measured output v . The block Δ represents structured or unstructured uncertainty in feedback with the nominal closed-loop system. Robust stability and performance are assessed with respect to this interconnection.

Consider the generalized plant interconnection shown in Figure 2.3. The block Δ represents uncertainty, which may account for parametric variations, unmodeled high-frequency dynamics, or exogenous disturbances. The controller K generates the control signal u from measured outputs v . The exogenous input w represents disturbances and reference inputs, while z denotes regulated performance outputs. The objective of robust control is to ensure internal stability of the closed-loop interconnection and bounded input–output gain from w to z for all admissible uncertainties Δ within a prescribed uncertainty set [4].

2.3.1 Norms

Norms provide quantitative measures of signal and operator amplification and form the mathematical foundation of robustness analysis. In control theory, stability and performance are often characterized in terms of induced gains between input and output signal spaces.

Vector and Signal Norms

For a vector $q \in \mathbb{R}^n$, the ℓ_1 -norm is defined as

$$\|q\|_1 = \sum_{i=1}^n |q_i|, \quad (2.33)$$

whereas the Euclidean norm (or ℓ_2 norm) is given by

$$\|q\|_2 = \sqrt{\sum_{i=1}^n |q_i|^2}. \quad (2.34)$$

Another commonly used norm is the peak norm which defines the largest amplification

$$\|q\|_\infty = \max_i |q_i|. \quad (2.35)$$

The norms in (2.33)–(2.35) can be extended to account for time dependent signals $q(t)$. In such cases the norms are evaluated over time horizon $t \in [0, \infty)$.

Matrix Norms

For a matrix $A \in \mathbb{R}^{m \times n}$, the induced (operator) norm corresponding to a given vector norm $\|\cdot\|$ is defined as

$$\|A\| = \sup_{x \neq 0} \frac{\|Ax\|}{\|x\|}, \quad (2.36)$$

where sup denotes the supremum, i.e., the least upper bound.

Of particular importance is the induced 2-norm,

$$\|A\|_2 = \bar{\sigma}(A), \quad (2.37)$$

where $\bar{\sigma}(A)$ denotes the largest singular value of A [4]. This norm characterizes the maximum Euclidean amplification of the linear mapping defined by A .

When the vector infinity norm is used, the induced matrix infinity norm becomes

$$\|A\|_\infty = \sup_{x \neq 0} \frac{\|Ax\|_\infty}{\|x\|_\infty} = \max_{1 \leq i \leq m} \sum_{j=1}^n |a_{ij}| = \max_{1 \leq i \leq m} \|A^{(i)}\|_1, \quad (2.38)$$

that is, the maximum ℓ_1 -norm of the rows of A .

Another frequently used matrix norm is the Frobenius norm,

$$\|A\|_F = \sqrt{\text{tr}(A^T A)}, \quad (2.39)$$

which equals the Euclidean norm, (2.34), of the singular value vector.

System Norms

For stable linear time-invariant (LTI) systems, robustness is characterized using induced system norms that describe input–output amplification. Let $G(s)$ be a stable transfer matrix with impulse response $g(t)$, where $s = j\omega$.

The induced \mathcal{L}_2 gain, also known as the \mathcal{H}_∞ norm, is defined as

$$\|G\|_{\mathcal{H}_\infty} = \sup_{u \in \mathcal{L}_2, u \neq 0} \frac{\|Gu\|_{\mathcal{L}_2}}{\|u\|_{\mathcal{L}_2}} = \sup_{\omega \in \mathbb{R}} \bar{\sigma}(G(j\omega)), \quad (2.40)$$

and represents the worst-case energy amplification from input to output.

Similarly, the induced \mathcal{L}_∞ gain is bounded using the \mathcal{L}_1 norm of the impulse response,

$$\|G\|_{\mathcal{L}_1} = \int_0^\infty \|g(t)\|_\infty dt, \quad (2.41)$$

which yields the input–output bound

$$\|Gu\|_{\mathcal{L}_\infty} \leq \|G\|_{\mathcal{L}_1} \|u\|_{\mathcal{L}_\infty}. \quad (2.42)$$

where u is the input. Thus, while the \mathcal{H}_∞ norm quantifies worst-case energy amplification and is evaluated in the frequency domain, the \mathcal{L}_1 norm is a time domain quantity that provides a bound on the worst-case peak output given a bounded input.

2.3.2 Small-Gain Theorem

A fundamental sufficient condition for robust stability of feedback interconnections is provided by the Small-Gain Theorem [11]. Consider the feedback interconnection of two stable causal operators G and Δ acting on \mathcal{L}_2 . If

$$\|G\|_\infty \|\Delta\|_\infty < 1, \quad (2.43)$$

then the feedback interconnection is internally stable.

In robust control applications, Δ typically represents normalized uncertainty satisfying $\|\Delta\|_\infty \leq 1$. In this case, a sufficient condition for robust stability reduces to

$$\|G\|_\infty < 1. \quad (2.44)$$

The small-gain condition formalizes the requirement that the loop amplification remain strictly below unity and constitutes the analytical basis for induced \mathcal{L}_2 control synthesis and modern robustness analysis [4].

2.4 \mathcal{L}_1 Adaptive Control

The \mathcal{L}_1 adaptive controller [2] establishes a systematic framework that combines fast parameter adaptation with robustness guarantees expressed through induced norm bounds. Conceptually, the architecture resembles the indirect MRAC formulation presented in Section 2.2.3, as it employs a state predictor to estimate parametric uncertainty, see Figure 2.4.

A defining feature of the \mathcal{L}_1 architecture is that the predictor is designed such that, under ideal parameter estimates $\hat{\Lambda} = \Lambda$ and $\hat{\Omega} = \Omega$, its dynamics coincide with the desired reference model. In other words, in the absence of estimation error, the predictor reproduces the reference model exactly. Consequently, the prediction error $\tilde{x} = \hat{x} - x$ vanishes and the adaptive loop does not influence the nominal closed-loop behaviour. This property provides a clear separation between nominal performance (governed by the reference model) and uncertainty compensation (handled by the adaptive mechanism).

In contrast to classical MRAC, the learning rates Γ_Λ and Γ_Θ in \mathcal{L}_1 control are deliberately chosen to be sufficiently fast. This enables rapid estimation of the matched uncertainty while ensuring that the control signal remains well behaved through appropriate filtering.

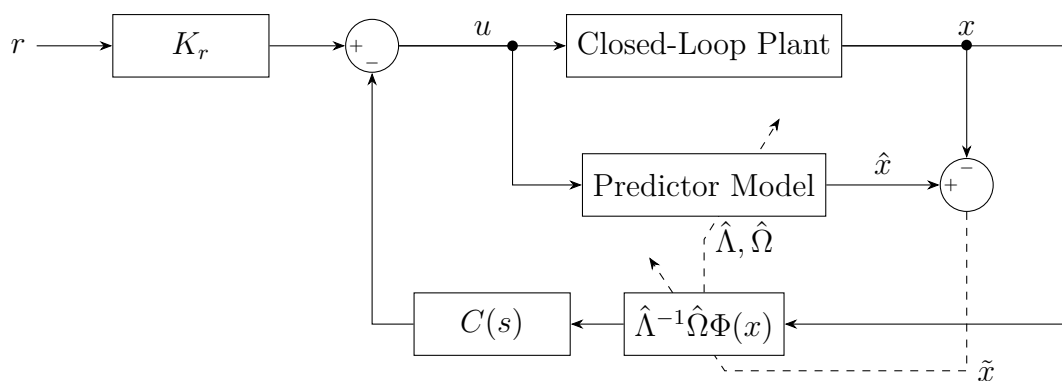


Figure 2.4: Block diagram of the \mathcal{L}_1 adaptive control architecture. The control input u consists of the nominal feedforward term $K_r^\top r$ and the adaptive component filtered through the strictly proper low-pass filter $C(s)$, i.e., $-C(s)\hat{\Lambda}^{-1}\hat{\Omega}^\top\Phi(x)$. The state predictor (2.22) generates the estimated states \hat{x} based on the measured plant state x from (2.21) and the control input. The prediction error $\tilde{x} = \hat{x} - x$ is used in the adaptive laws to update the uncertainty estimates $\hat{\Omega}$ and $\hat{\Lambda}$.

As illustrated in Figure 2.4, a strictly proper low-pass filter [12] is introduced in the control channel to suppress high-frequency oscillations induced by fast adaptation. This filtering is essential for achieving a decoupling between adaptation and robustness, thereby preserving transient performance guarantees while allowing high adaptation gains.

However, the introduced low-pass filter fundamentally alters the control input perceived by the plant and prediction error. As a result, the adaptive laws no longer has access to the true excitation required for parameter convergence. Thus, even if the unfiltered regressor in (2.45) satisfies the PE condition in Section 2.2.4, the filtered may not.

Therefore, while boundedness and transient performance are preserved, the low-pass filter effectively negates PE in the adaptive loop, preventing convergence of the parameter estimates to their true values in general.

The applied control input is given by

$$u = K_r^\top r - C(s)\hat{\Lambda}^{-1}\hat{\Omega}^\top\Phi(x), \quad (2.45)$$

where $C(s)$ denotes the low-pass filter

$$C(s) = \frac{\omega_c}{s + \omega_c}, \quad (2.46)$$

with $\omega_c > 0$ representing the cutoff frequency [13].

The filter bandwidth is selected such that a small-gain type condition is satisfied in the \mathcal{L}_∞ framework [2]. Specifically, defining

$$H(s) = (sI - A_m)^{-1}B, \quad (2.47)$$

which is the transfer function of (2.21), the design condition

$$\|H(s)(1 - C(s))\|_{\mathcal{L}_1} \|\Omega\|_\infty < 1 \quad (2.48)$$

guarantees bounded-input bounded-output stability of the adaptive interconnection. Here, $\|\Omega\|_\infty$ denotes the worst-case bound on the uncertain parameter vector, while $\|H(s)(1 - C(s))\|_{\mathcal{L}_1}$ represents the induced $\mathcal{L}_\infty \rightarrow \mathcal{L}_\infty$ gain of the filtered uncertainty channel. Condition (2.48) can therefore be interpreted as a small-gain constraint that guarantees robustness of the adaptive closed-loop system independently of the adaptation rate.

In particular, robustness is improved as the product $\|H(s)(1 - C(s))\|_{\mathcal{L}_1} \|\Omega\|_\infty$ decreases, i.e., when the uncertainty channel is sufficiently attenuated. However, this attenuation is achieved through increased filtering, which reduces the effective control bandwidth and consequently degrades tracking performance.

Thus, tuning the filter introduces an inherent trade-off between robustness to uncertainty and closed-loop tracking performance.

2.5 Projection Operator

To further guarantee boundedness of the parameter estimates and to prevent drift under high adaptation gains, a projection operator inspired by the bounded parameter projection techniques commonly used in adaptive control [2] is incorporated into the adaptive laws, i.e. (2.20) and (2.30). The projection ensures that the estimates remain within a prescribed compact and convex set, thereby preserving robustness while allowing fast adaptation independently of the adaptation gain Γ .

2.5.1 Convex Set and Smooth Convex Function

The projection operator enforces membership in the smooth convex set

$$\mathcal{D} = \{\theta \in \mathbb{R}^n \mid f(\theta) \leq c\}, \quad 0 \leq c \leq 1, \quad (2.49)$$

where $f(\cdot) : \mathbb{R}^n \rightarrow \mathbb{R}$ is the smooth convex function

$$f(\theta) = \frac{(\epsilon + 1)\theta^\top \theta - \theta_{\max}^2}{\epsilon \theta_{\max}^2}. \quad (2.50)$$

Here, $\theta_{\max} > 0$ is the prescribed bound on $\|\theta\|$ and $\epsilon > 0$ is the projection tolerance that defines the thickness of the boundary layer within which the projection

gradually engages. A larger ϵ results in a wider boundary layer and smoother projection behavior, while a smaller ϵ results in a thinner boundary layer closer to a hard constraint. The gradient of f is

$$\nabla f(\theta) = \frac{2(\epsilon + 1)\theta}{\epsilon\theta_{\max}^2}, \quad (2.51)$$

which is proportional to θ and therefore points radially outward. Since $\nabla f \propto \theta$, the unit normal to the boundary simplifies to

$$\frac{\nabla f}{\|\nabla f\|_2} = \frac{\theta}{\|\theta\|_2}. \quad (2.52)$$

2.5.2 Projection Operator Definition

The projection operator is defined as

$$\text{Proj}(\theta, y) = \begin{cases} y, & \text{if } f(\theta) < 0, \\ y, & \text{if } f(\theta) \geq 0 \text{ and } \nabla f^\top y \leq 0, \\ y - \frac{\nabla f \nabla f^\top}{\|\nabla f\|_2^2} y f(\theta), & \text{if } f(\theta) \geq 0 \text{ and } \nabla f^\top y > 0. \end{cases} \quad (2.53)$$

The three cases admit the following geometric interpretation. In the first case, θ lies strictly inside \mathcal{D} and the update y is applied freely. In the second case, θ is in or beyond the boundary layer but the update points inward, so no modification is needed. In the third case, θ is in or beyond the boundary layer and the update points outward where the projection removes a fraction $f(\theta)$ of the outward component of y in the direction of ∇f . Since $\nabla f \propto \theta$, this simplifies to

$$\text{Proj}(\theta, y) = y - \frac{\theta\theta^\top}{\|\theta\|_2^2} f(\theta) y, \quad f(\theta) \geq 0, \nabla f^\top y > 0, \quad (2.54)$$

which shows that only the component of y in the direction of θ is modified, and the modification is proportional to how deep into the boundary layer θ has penetrated. At $f(\theta) = 0$ no modification occurs, while at $f(\theta) = 1$, i.e. $\|\theta\|_2 = \theta_{\max}$, the outward component is fully removed.

2.5.3 Key Properties

The projection operator satisfies two properties that are essential for the stability analysis. First, it guarantees boundedness of the parameter estimates for all time,

$$\|\hat{\theta}(t)\|_2 \leq \theta_{\max}, \quad \forall t \geq 0, \quad (2.55)$$

provided $\|\hat{\theta}(0)\|_2 \leq \theta_{\max}$. This follows directly from the fact that at $\|\theta\|_2 = \theta_{\max}$, $f(\theta) = 1$ and the full outward component of the update is removed, preventing further growth. Second, the projection satisfies the Lyapunov compatibility condition

$$\tilde{\theta}^\top (\text{Proj}(\hat{\theta}, y) - y) \leq 0, \quad (2.56)$$

where $\tilde{\theta} = \hat{\theta} - \theta$ is the parameter estimation error. This inequality ensures that the projection does not interfere with the convergence properties established in the Lyapunov analysis of the adaptive laws, i.e. the projection can only improve or maintain the Lyapunov decrease, never worsen it.

2.6 Linear Quadratic Regulator

The Linear Quadratic Regulator (LQR) is a classical optimal state-feedback controller for LTI systems. Consider the system

$$\begin{aligned}\dot{x} &= Ax + Bu, \\ y &= Cx + Du,\end{aligned}\tag{2.57}$$

where $x \in \mathbb{R}^n$ is the state vector, $u \in \mathbb{R}^m$ is the control input, and $y \in \mathbb{R}^p$ is the system output. The matrices $A \in \mathbb{R}^{n \times n}$ and $B \in \mathbb{R}^{n \times m}$ represent the system dynamics and input matrix, respectively, while $C \in \mathbb{R}^{p \times n}$ and $D \in \mathbb{R}^{p \times m}$ define the measured output.

The LQR aims to determine a state-feedback control law

$$u = -K_x x,\tag{2.58}$$

that minimizes the infinite-horizon quadratic cost function

$$J = \int_0^\infty (x^\top Q_x x + u^\top R u + 2x^\top Q_{xu} u) dt,\tag{2.59}$$

where $Q_x \in \mathbb{R}^{n \times n}$ and $R \in \mathbb{R}^{m \times m}$ are symmetric weighting matrices, $Q_x = Q_x^\top \geq 0$ and $R = R^\top > 0$, chosen to penalize deviations in states and control effort, respectively. The cost matrix $Q_{xu} \in \mathbb{R}^{n \times m}$ introduces cross-weighting between the states and the control input, allowing the cost to penalize their interaction in addition to the individual state and control magnitudes [6].

The optimal gain K_x is obtained by solving the continuous-time Algebraic Riccati Equation (CARE)

$$A^\top S + SA - (SB + Q_{xu}) R^{-1} (SB + Q_{xu})^\top + Q_x = 0,\tag{2.60}$$

where $S = S^\top \geq 0$ is the unique stabilizing solution ensuring asymptotic stability of the closed-loop system. The resulting optimal feedback gain is

$$K_x = R^{-1} (SB + Q_{xu})^\top.\tag{2.61}$$

The closed-loop dynamics under LQR feedback are then

$$\dot{x} = (A - BK_x)x,\tag{2.62}$$

which are guaranteed to be asymptotically stable for stabilizable matrices (A, B) and detectable matrices $(A, Q_x^{1/2})$ [6].

For reference tracking of constant signals, the control law in (2.58) can be extended to

$$u = -K_x x + K_r r,\tag{2.63}$$

where $r \in \mathbb{R}^p$ is the reference input, and K_r is a feedforward gain chosen to eliminate steady-state error. Assuming $(A - BK_x)$ is nonsingular and the right hand side is invertible, K_r is computed from

$$K_r = - [C(A - BK_x)^{-1}B + D]^{-1}. \quad (2.64)$$

This formulation provides a systematic method to balance state regulation against control effort, with the weighting matrices Q and R determining the trade-off. The LQ controller, summarized in Figure 2.5, is widely used in both theory and practice due to its optimality, simplicity, and inherent robustness to moderate model uncertainties [6].




Figure 2.5: Block diagram of the LQ state-feedback controller with reference feedforward gain. The control input is given by $u = K_r r - K_x x$, where K_x is the optimal state-feedback gain obtained by minimizing a quadratic cost function, and K_r is the reference feedforward gain ensuring steady-state tracking of the reference signal r . The states x from the plant (2.57) is fed back through the gain K_x to regulate the system states and achieve optimal performance in the sense of the chosen quadratic cost functional.

2.6.1 Linear Parameter-Varying Systems

Systems whose dynamics depend on measurable, time-varying parameters are commonly modeled as Linear Parameter-Varying (LPV) systems. An LPV system is described by

$$\begin{aligned} \dot{x} &= A(\rho)x + B(\rho)u, \\ y &= C(\rho)x + D(\rho)u, \end{aligned} \quad (2.65)$$

where $\rho \in \mathcal{P} \subset \mathbb{R}^s$ is a measurable scheduling parameter belonging to a known bounded set \mathcal{P} , typically assumed to be a compact polytope defined by

$$\rho \in [\rho_{\min}, \rho_{\max}].$$

The parameter ρ may represent operating conditions such as velocity, altitude, or other nonlinear effects, allowing certain nonlinear systems to be embedded into an LPV framework.

Stability analysis of LPV systems is commonly performed using Lyapunov methods. A sufficient condition for quadratic stability is the existence of a constant symmetric matrix $P = P^\top > 0$ such that

$$A(\rho)^\top P + PA(\rho) < 0, \quad \forall \rho \in \mathcal{P}, \quad (2.66)$$

which guarantees exponential stability uniformly over the parameter set [4]. This condition is often referred to as the existence of a common quadratic Lyapunov function.

If no constant matrix P satisfies the above inequality for all ρ , stability can be analyzed using a parameter-dependent Lyapunov function $V(x, \rho) = x^\top P(\rho)x$, leading to the condition

$$\frac{\partial P(\rho)}{\partial \rho} \dot{\rho} + A(\rho)^\top P(\rho) + P(\rho)A(\rho) < 0, \quad (2.67)$$

which explicitly accounts for the rate of variation of the scheduling parameter.

2.6.2 Gain Scheduling

When the plant is described by an LPV model, controller design is typically performed using gain scheduling. In this approach, a family of controllers

$$u = -K_x(\rho)x + K_r(\rho)r$$

is constructed, where the feedback, $K_x(\rho)$, and feedforward, $K_r(\rho)$, gains depend explicitly on the scheduling parameter. The controller gains are either interpolated between locally designed linear controllers corresponding to fixed operating points, or synthesized directly within the LPV [14].

Gain scheduling provides a practical means of extending linear control techniques to systems operating over wide ranges of conditions. However, stability of the overall closed-loop system requires careful analysis to ensure that the scheduled controller preserves stability for all admissible parameter trajectories.

2.7 Parameter Summary

A summary to clarify which parameters are time dependent and which are signals is presented in Table 2.1.

Table 2.1: Parameters and their dependencies

Notation	Explicit argument	Description
x	$x(t)$	Measured states
\hat{x}	$\hat{x}(t)$	Estimated states
\tilde{x}	$\tilde{x}(t)$	Prediction error
x_m	$x_m(t)$	Reference states
e	$e(t)$	Tracking error
r	$r(t)$	Reference signal
u	$u(t)$	Control input
A	A	LTI plant matrix
B	B	LTI input matrix
C	C	LTI output matrix
D	D	LTI feed-through matrix
A_m	A_m	LTI Reference plant matrix
B_m	B_m	LTI Reference plant matrix
A_Δ	$A_\Delta(t)$	Plant matrix deviations
B_Δ	$B_\Delta(t)$	Input matrix deviations
$V(x)$	$V(x, t)$	Lyapunov candidate function
Q	Q	LTI system penalty
P	P	LTI Lyapunov solution
Λ	$\Lambda(t)$	Input uncertainty
$\hat{\Lambda}$	$\hat{\Lambda}(t)$	Estimated input uncertainty
$\tilde{\Lambda}$	$\tilde{\Lambda}(t)$	Input uncertainty estimation error
$\Delta(x)$	$\Delta(x, t)$	Model uncertainty
Θ	$\Theta(t)$	Unknown uncertainty coefficients
$\hat{\Theta}$	$\hat{\Theta}(t)$	Estimated uncertainty coefficients
$\tilde{\Theta}$	$\tilde{\Theta}(t)$	Uncertainty coefficients estimation error
Ω	$\Omega(t) = \Lambda(t)\Theta(t)$	Scaled uncertainty weights
$\hat{\Omega}$	$\hat{\Omega}(t)$	Estimated scaled uncertainty weights
$\tilde{\Omega}$	$\tilde{\Omega}(t)$	Scaled uncertainty weights error
$\Phi(x)$	$\Phi(x, t)$	Lipschitz-continuous basis functions
K_x	K_x	Reference feedback gain
\hat{K}_x	$\hat{K}_x(t)$	Estimated feedback gain
\tilde{K}_x	$\tilde{K}_x(t)$	Feedback gain estimation error
K_r	K_r	Reference feedforward gain
\hat{K}_r	$\hat{K}_r(t)$	Estimated feedforward gain
\tilde{K}_r	$\tilde{K}_r(t)$	Feedforward gain estimation error
ρ	$\rho(t)$	Time varying parameter
$A(\rho)$	$A(\rho, t)$	LPV plant matrix
$B(\rho)$	$B(\rho, t)$	LPV input matrix
$C(\rho)$	$C(\rho, t)$	LPV output matrix
$D(\rho)$	$D(\rho, t)$	LPV feed-through matrix
$K_x(\rho)$	$K_x(\rho, t)$	Gain scheduled feedback gain
$K_r(\rho)$	$K_r(\rho, t)$	Gain scheduled feedforward gain

3

Methods

This chapter presents the methodology used in designing a robust adaptive controller.

3.1 System Description

This section gives an overview of the pre-implemented simulation model.

3.1.1 Vehicle Structure

The simulated vehicle is assumed to be rotationally symmetric about its longitudinal axis. To steer the vehicle, four control surfaces, arranged in an four-fin X-configuration, are located in the rear. Figure 3.1 shows an illustration of the vehicle where v is the velocity vector. The attack α and side-slip angle β are the angle deviations between the x and y coordinate axis to the velocity vector v , respectively.




Figure 3.1: Illustration of Aerial Vehicle. Figure is obtained from [1].

For simplicity, the controllers implemented in this report are designed for traditional aircraft, which use aileron, elevator, and rudder as control surfaces. The control inputs given by the control law are thus recomputed from the mentioned configuration into the X-configuration using:

$$u_{\delta_{1,2,3,4}} = \begin{bmatrix} 1 & -1 & 1 \\ -1 & 1 & 1 \\ 1 & 1 & -1 \\ -1 & -1 & -1 \end{bmatrix} u_{\delta_{\text{AER}}} \quad (3.1)$$

where $u_{\delta_{1,2,3,4}}$ is a vector of the recomputed control inputs for the X-configuration control surfaces, while $u_{\delta_{\text{AER}}}$ represents the vector of raw control inputs given by the controller.

The mathematical representation of the model can be summarized as

$$\dot{x} = f(x, u_{\delta_{\text{AER}}}) \quad (3.2)$$

where $f(\cdot)$ is the non-linear model.

3.1.2 Simulation Model

To be able to apply the control architectures presented in Chapter 2, the non-linear model in (3.2) is split into three sub-models controlling roll, pitch, and yaw, respectively, where each model is derived assuming the other two channels are not excited. Thus, all cross-coupling are omitted, which is generally valid when assuming motion only occurs predominantly in a single channel. To balance model accuracy against local linearization errors, the nonlinear model (3.2) is linearized at a set of predefined operating points within the flight envelope, as illustrated in Figure 3.2, rather than about the instantaneous operating point.

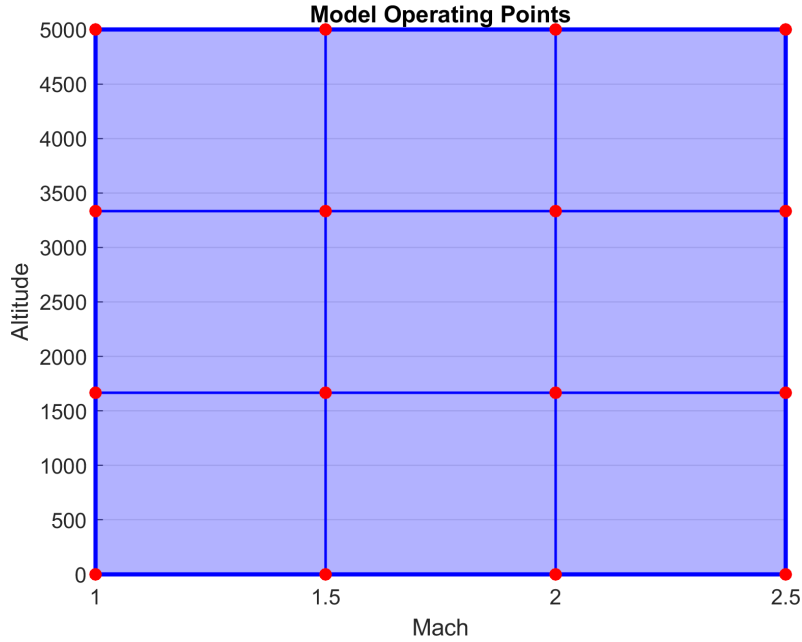


Figure 3.2: Polytope of operating-points for model linearization.

The resulting LPV models are summarized as:

$$\begin{aligned} \dot{x}_\varphi &= A_\varphi(v, h)x_\varphi + B_\varphi(v, h)u_{\delta_A}, \\ \dot{x}_\vartheta &= A_\vartheta(v, h)x_\vartheta + B_\vartheta(v, h)u_{\delta_E}, \\ \dot{x}_\psi &= A_\psi(v, h)x_\psi + B_\psi(v, h)u_{\delta_R}, \end{aligned} \quad (3.3)$$

where φ , ϑ , and ψ are the three Euler angles representing roll, pitch, and yaw, respectively, and v and h are the velocity and the altitude, respective. As part of the interpolation process, the velocity is recomputed from m/s to Mach using

$$M = \frac{\|v\|_2}{\sqrt{\gamma RT}} \quad (3.4)$$

where γ , R , and T are the ratio of specific heat, gas constant, and absolute temperature, respectively [15].

Thus, the model for roll in (3.3) is controlled by the aileron u_{δ_A} , the model for pitch is controlled by the elevator u_{δ_E} , while the model for yaw is controlled by the rudder u_{δ_R} .

It is further assumed that all states are measured, enabling full-state feedback. Table 3.1 summarizes all measured states, where x_φ , x_ϑ , and x_ψ contain the following states

$$x_\varphi = \begin{bmatrix} \varphi \\ \dot{\varphi} \\ \int (\varphi_c - \varphi) dt \end{bmatrix}, \quad x_\vartheta = \begin{bmatrix} \alpha \\ \dot{\vartheta} \\ \int (a_{z_c} - a_z) dt \end{bmatrix}, \quad x_\psi = \begin{bmatrix} \beta \\ \dot{\psi} \\ \int (a_{y_c} - a_y) dt \end{bmatrix}, \quad (3.5)$$

where φ_c , a_{z_c} , and a_{y_c} are the commanded roll angle, acceleration along z-axis, and acceleration along y-axis, respectively, and are given by the guidance module, further described in Section 3.1.4, to obtain desired control surface deflection.

Table 3.1: Measured States

Notation	States	Unit
h	Altitude	m
φ	Roll angle	m
α	Attack angle	rad
β	Side-slip angle	rad
v_x	Velocity along x -axis	m/s
v_y	Velocity along y -axis	m/s
v_z	Velocity along z -axis	m/s
a_x	Acceleration along x -axis	m/s ²
a_y	Acceleration along y -axis	m/s ²
a_z	Acceleration along z -axis	m/s ²
$\dot{\varphi}$	Roll angular velocity	rad/s ²
$\dot{\vartheta}$	Pitch angular velocity	rad/s ²
$\dot{\psi}$	Yaw angular velocity	rad/s ²
δ_A	Aileron control surface deflection	rad
δ_E	Elevator control surface deflection	rad
δ_R	Rudder control surface deflection	rad

With both velocity v and altitude h available, an arbitrary operating point within the polytope shown in Figure 3.2 is obtained during simulation by first identifying the corresponding cell (square) of the partition in which the state lies. Within each cell, system matrices are computed using bilinear interpolation of the surrounding corner points.

The interpolation weights are defined as

$$\begin{aligned}
 h_h &= \frac{h - h_{\min}}{h_{\max} - h_{\min}}, & h_l &= 1 - h_h, \\
 v_h &= \frac{v - v_{\min}}{v_{\max} - v_{\min}}, & v_l &= 1 - v_h, \\
 A(v, h) &= (A(v_{\min}, h_{\min})h_l - A(v_{\min}, h_{\max})h_h)v_l + \\
 & \quad (A(v_{\max}, h_{\min})h_l - A(v_{\max}, h_{\max})h_h)v_h
 \end{aligned} \tag{3.6}$$

where h_{\min} , h_{\max} and v_{\min} , v_{\max} denote the altitude and velocity bounds of the corresponding cell and $A(\cdot)$ represents the velocity and altitude dependent model matrices in (3.3). The velocity is expressed in Mach number to ensure consistency with the operating points in Figure 3.2. The resulting interpolation yields a smooth approximation of the system matrices over the partitioned operating domain and may result in injected model errors depending on the nonlinearities of the true model in (3.2).

3.1.3 Baseline Controller

With the system model split into three sub-models, three nominal controllers are designed. The nominal controllers are designed using the LQ framework, presented in Section 2.6, and are tuned for four operating-points within the polytope illustrated in Figure 3.2, specifically the points shown in Figure 3.3.




Figure 3.3: Polytope of operating-points for gain scheduling.

The gain scheduled control laws can therefore be written as

$$\begin{aligned}
 u_{\delta_A} &= -K_{x_\varphi}(v, h)x_\varphi + K_{r_\varphi}(v, h)r_\varphi, \\
 u_{\delta_E} &= -K_{x_\vartheta}(v, h)x_\vartheta + K_{r_\vartheta}(v, h)r_\vartheta, \\
 u_{\delta_R} &= -K_{x_\psi}(v, h)x_\psi + K_{r_\psi}(v, h)r_\psi,
 \end{aligned} \tag{3.7}$$

To obtain gains at arbitrary operating conditions, bilinear interpolation between the four operating points is employed. The interpolated gain matrix is computed as

$$\begin{aligned}
h_h &= \frac{h - h_{\min}}{h_{\max} - h_{\min}}, & h_l &= 1 - h_h, \\
v_{\max} &= v_{\max,1} + h_h(v_{\max,2} - v_{\max,1}), \\
v_h &= \frac{\|v\|_2 - v_{\min}}{v_{\max} - v_{\min}}, & v_l &= 1 - v_h \\
K(v, h) &= (K(v_{\min}, h_{\min})h_l + K(v_{\min}, h_{\max})h_h)v_l \\
&\quad + (K(v_{\max,1}, h_{\min})h_l + K(v_{\max,2}, h_{\max})h_h)v_h,
\end{aligned} \tag{3.8}$$

where $v_{\max,1}$ and $v_{\max,2}$ are the maximum velocity at low and high altitude, respectively, and $K(\cdot)$ denotes the altitude and velocity dependent controller gain matrices. Thus, similar to the models in (2.57), with interpolation being used, error may be injected in the model dynamics.

Since the models in (3.3) are derived under the assumption of decoupled rotational dynamics, nonlinear cross-coupling effects may arise when multiple channels are simultaneously excited. These effects introduce modeling errors that are not captured by the linearized sub-models. To mitigate the resulting performance degradation, integral states are introduced as defined in (3.5). The integral action improves robustness against cross-coupling effects and compensates for modeling uncertainties and nonlinearities.

3.1.4 Guidance Module

The commanded accelerations a_{z_c} and a_{y_c} , along with the roll angle φ_c , are generated by a pre-implemented guidance module. This module derives the commanded signals based on a series of defined way-points that the vehicle is required to follow.

The way-points are specified in the sequential order in which the vehicle must navigate them. Initially, the guidance module directs the vehicle towards the target way-point. As the vehicle advances, the module utilises a set of predefined conditions to determine the optimal moment to transition to the subsequent way-point, thereby ensuring a smooth and continuous trajectory.

As the vehicle approaches the current point, the module assesses whether the vehicle is within a specified turn circle distance. This distance is calculated based on the ideal turn radius between consecutive legs (the path from one way-point to the next) and the angular deviation between the current and next way-points. If the vehicle is within this turn circle, it initiates a turn towards the next point, preparing for the transition.

Additionally, the module verifies whether the vehicle has passed the current point by confirming that it is positioned ahead of the way-point relative to its velocity vector. This ensures that the vehicle has adequately progressed along the current leg before proceeding to the next point.

Figure 3.4 provides an illustration of how the guidance module directs the vehicle to follow a sequence of three way-points. The target way-point is considered passed once the vehicle is located ahead of it, facilitating a seamless transition to the next way-point in the sequence.




Figure 3.4: Example of guidance module commanding a vehicle

3.2 Design of Robust Adaptive Control

As discussed in Section 2.4, the \mathcal{L}_1 adaptive control architecture relies on the assumption that the state predictor reproduces the desired reference model dynamics under ideal parameter estimates $\hat{\Lambda}$ and $\hat{\Omega}$. A fundamental requirement for guaranteeing transient performance is the existence of a strictly proper filter $C(s)$ satisfying the small-gain condition in (2.48).

However, if the selected reference model exhibits excessively slow dynamics, it may be impossible to construct a filter $C(s)$ that satisfies (2.48). In such a case, the predictor can no longer mimic the reference model, and the nominal \mathcal{L}_1 framework does not provide the intended tracking guarantees. This limitation motivates a modification of the architecture to preserve robustness and bounded tracking performance despite the mismatch.

To address this issue, the controller derived in [2] is extended into a cascade-like configuration [16]. The outer loop is formulated as a direct MRAC (see Section 2.2.2) to ensure bounded tracking error with respect to the reference model. The inner loop retains the structure of an \mathcal{L}_1 adaptive controller but employs a sufficiently fast predictor to guarantee that (2.48) holds independently of the outer-loop dynamics. This architecture is conceptually related to the composite MRAC structure introduced in [17] and will hereafter be referred to as the \mathcal{L}_1 CMRAC.

Similarly to the direct MRAC formulation, the plant is modeled according to (2.6). The applied control input combines the adaptive state-feedback and feedforward terms from the direct MRAC (2.15) with the filtered \mathcal{L}_1 compensation from (2.45):

$$u = \hat{K}_r^\top r - \hat{K}_x^\top x - C(s)\hat{\Lambda}^{-1}\hat{\Omega}^\top \Phi(x). \quad (3.9)$$

3.2.1 Inner Loop

The predictor in (2.22) is modified to

$$\dot{\hat{x}} = A_l \hat{x} + B \left(\hat{\Lambda} u + \hat{\Omega}^\top \Phi(x) \right), \quad A_l = A - BL, \quad (3.10)$$

where L is a stabilizing feedback gain selected such that A_l is Hurwitz with sufficiently fast eigenvalues to ensure that a modified small-gain condition (2.48) is satisfied. The modification is described in Section 3.4.

Using (3.10) together with (2.6), the prediction error dynamics becomes

$$\begin{aligned}\dot{\tilde{x}} &= A_l \tilde{x} + B \left(\hat{\Lambda} u + \hat{\Omega}^\top \Phi \right) - \left(A x + B \left(\Lambda u + \Omega^\top \Phi \right) \right) \\ &= A_l \tilde{x} + B \left(\tilde{\Lambda} u + \tilde{\Omega}^\top \Phi(x) \right),\end{aligned}\quad (3.11)$$

where $\tilde{\Lambda} = \hat{\Lambda} - \Lambda$ and $\tilde{\Omega} = \hat{\Omega} - \Omega$.

Since $A_l = A - BL$, the term Lx enters the state dynamics through the same input channel as the matched uncertainty. Consequently, it can be absorbed into the lumped uncertainty representation $\tilde{\Omega}^\top \Phi(x)$. This modeling choice preserves the matched uncertainty structure but introduces a bias in the parameter estimates. In particular, even under PE (see Section 2.2.4), convergence of $\hat{\Omega}$ to the true parameter vector is not guaranteed; instead, convergence occurs to a neighborhood determined by L . Nevertheless, boundedness of the estimation errors is retained.

Boundedness of the inner-loop signals is established using the Lyapunov function in (2.27), where $P_l = P_l^\top > 0$ satisfies (2.5) for $A = A_l$. The time derivative is given by

$$\dot{V} \leq -\tilde{x}^\top Q_l \tilde{x} + 2\text{tr} \left(\tilde{\Lambda}^\top \left[\Gamma_\Lambda^{-1} \dot{\hat{\Lambda}} + u \tilde{x}^\top P_l B \right] \right) + 2\text{tr} \left(\tilde{\Omega}^\top \left[\Gamma_\Omega^{-1} \dot{\hat{\Omega}} + \Phi \tilde{x}^\top P_l B \right] \right), \quad (3.12)$$

where Q_l is defined through (2.5).

By selecting the update laws according to (2.30), the cross terms are cancelled, and \dot{V} becomes negative semi-definite. Hence, \tilde{x} , $\tilde{\Lambda}$, and $\tilde{\Omega}$ remain bounded for all time.

3.2.2 Outer Loop

The outer loop follows the direct MRAC derivation presented in Section 2.2.2. Substituting (3.9) into (2.16) and using (2.14) yields the tracking error dynamics

$$\dot{e} = A_m e + B \Lambda \left(\tilde{K}_r^\top r - \tilde{K}_x^\top x + \left[\Theta^\top - C(s) \hat{\Lambda}^{-1} \hat{\Omega}^\top \right] \Phi(x) \right). \quad (3.13)$$

Consider the Lyapunov candidate

$$V(e, \tilde{K}_r, \tilde{K}_x) = e^\top P e + \text{tr} \left(\left[\tilde{K}_r^\top \Gamma_r^{-1} \tilde{K}_r + \tilde{K}_x^\top \Gamma_x^{-1} \tilde{K}_x \right] \Lambda \right), \quad (3.14)$$

where P satisfies (2.5) for $A = A_m$.

Its time derivative becomes

$$\begin{aligned}\dot{V} &\leq -e^\top Q_m e + 2\text{tr} \left(\tilde{K}_r^\top \left[\Gamma_r^{-1} \dot{\tilde{K}}_r + r e^\top P_m B \right] \Lambda \right) \\ &\quad + 2\text{tr} \left(\tilde{K}_x^\top \left[\Gamma_x^{-1} \dot{\tilde{K}}_x - x e^\top P_m B \right] \Lambda \right) \\ &\quad - 2e^\top P_m B \Lambda \left(\Theta^\top - C(s) \hat{\Lambda}^{-1} \hat{\Omega}^\top \right) \Phi(x).\end{aligned}\quad (3.15)$$

In contrast to the standard direct MRAC case, the update laws (2.20) do not cancel all cross terms due to the presence of the filtered adaptive component. Nevertheless, the inner-loop analysis guarantees boundedness of $\tilde{\Lambda}$ and $\tilde{\Omega}$ via (3.12). Consequently, the residual term in (3.15) remains bounded.

Applying the composite Lyapunov framework in [18], it follows that all closed-loop signals are bounded and that the tracking error e converges to a bounded neighborhood of the origin. The size of this neighborhood is determined by the inner-loop estimation accuracy and the bandwidth of the filter $C(s)$.

3.2.3 Extended Projection Operator

With the adaptive laws applied in discrete time steps, the projection operator described in Section 2.5 is extended to the matrix case and augmented with a hard constraint safeguard to handle large discrete update steps that could otherwise violate the projection set in a single iteration.

Specifically, the proposed operator $\text{Proj}(\cdot)$ enforces, in addition to (2.55), a bound on the induced infinity norm of the parameter matrix by constraining each row vector $\theta^{(i)}$ to lie inside an ℓ_1 -ball of prescribed radius $\theta_{\max} > 0$. Recalling that the induced infinity norm of a matrix equals the maximum ℓ_1 -norm of its rows (2.38), this construction guarantees

$$\|\theta\|_{\infty} \leq \theta_{\max}. \quad (3.16)$$

For a nominal update direction y , the projection is defined row-wise as

$$\text{Proj}(\theta_k, y_k) = \begin{cases} \theta_k^{(i)} + y_k^{(i)}, & \text{if } \|\theta_k^{(i)} + y_k^{(i)}\|_1 \leq \theta_{\max}, \\ \frac{\theta_{\max}}{\|\theta_k^{(i)} + y_k^{(i)}\|_1} (\theta_k^{(i)} + y_k^{(i)}), & \text{otherwise.} \end{cases} \quad (3.17)$$

Here, the ℓ_1 -ball of radius θ_{\max} is defined as

$$\mathcal{S}(\theta_{\max}) = \{\theta_k \in \mathbb{R}^n : \|\theta_k\|_1 \leq \theta_{\max}\},$$

which is a convex, compact set, where the index k notates the current step along the time series t . If the tentative update remains inside this set, it is applied unchanged. Otherwise, it is radially scaled back onto the boundary of the ℓ_1 -ball, ensuring that the constraint is satisfied. The concept is visualized in Figure 3.5 where the white line shows the unaltered step, whereas the red line illustrates the correct update, and $\theta_{k+1}^{(i)}$ represents the updated row. The filled in gray area represents the feasible region within the convex set $\mathcal{S}(\theta_{\max})$.

This construction guarantees bounded parameter estimates, i.e. (3.16), while preserving continuity of the adaptive law and maintaining the stability properties established in the Lyapunov analysis.

3.3 Design of Reference Model

To ensure bounded tracking error while enabling a fast predictor in the inner loop, the reference model is designed to represent the nominal closed-loop behavior of the plant. Specifically, the gain scheduled controllers described in Section 3.1.3 are applied to the LPV systems derived in Section 3.1.2, . Thus, the reference model become:

$$\begin{aligned} A_m(v, h) &= A(v, h) - B(v, h)K_x(v, h) \\ B_m(v, h) &= B(v, h)K_r(v, h) \end{aligned} \quad (3.18)$$

where $A(\cdot)$ and $B(\cdot)$ are interpolated using (3.6) between the operating points illustrated in Figure 3.2, while $K_x(\cdot)$ and $K_r(\cdot)$ are interpolated using (3.8) between the points shown in Figure 3.3. This framework is applied to all control channels.




Figure 3.5: 2D illustration of row-wise projection of the parameter update onto the ℓ_1 -ball in the extended projection operator. The red line represents the correct update step

3.4 Design of Prediction Model

As mentioned in Section 2.4, a sufficiently small \mathcal{L}_1 norm improves robustness at the cost of closed-loop tracking performance. However, due to the model being updated each time instance through interpolation, a model error is introduced, which is amplified by the predictor dynamics. Thus, the predictor must be sufficiently fast to ensure (2.48) holds, while not causing unbounded prediction error \tilde{x} by amplified model error. This is ensured by modifying the small gain condition in (2.48) for the \mathcal{L}_1 CMRAC.

Starting from the prediction error (3.11), the applied control law (3.9) is explicitly written yielding:

$$\dot{\tilde{x}} = A_l(v, h)\tilde{x} + B(v, h) \left([1 - C(s) - \Lambda C(s)\hat{\Lambda}^{-1}] \hat{\Omega}^\top \Phi(x) + \text{other terms} \right), \quad (3.19)$$

where “other terms” include contributions from the outer loop control and the true plant uncertainties. These other terms are either stable or handled by the Lyapunov-based adaptive laws and are omitted for the purpose of the \mathcal{L}_1 analysis.

The predictor matrix $A_l(\cdot)$ is computed similarly to $A_m(\cdot)$ in (3.18), according to

$$A_l(v, h) = A(v, h) - B(v, h)L(v, h), \quad (3.20)$$

where $L(\cdot)$ denotes the scheduled predictor gain, selected to ensure sufficiently fast predictor dynamics. The gain $L(\cdot)$, together with the low-pass filter $C(s)$, is scheduled using the same approach as the baseline controller described in Section 3.1.3. The term $1 - C(s) - \Lambda C(s)\hat{\Lambda}^{-1}$ represents the residual uncertainty, which is the high-frequency content of the inner loop adaptive signal in (3.9) that is attenuated by the low-pass filter and therefore never reaches the plant, but still drives the prediction error dynamics. The residual uncertainty is therefore the component of the adaptive signal that could destabilize the system if not properly bounded. Since

the uncertain parameters Ω are bounded in amplitude, the \mathcal{L}_1 norm (2.41) provides the natural worst-case bound on the peak effect of the residual uncertainty on the prediction error dynamics, as established in Section 2.3.1. The stability condition is therefore expressed in terms of the transfer functions $H(s)$ and $C(s)$, which are naturally defined in the frequency domain, allowing the \mathcal{L}_1 norm to be evaluated numerically.

Treating $\hat{\Omega}^\top \Phi(x)$ as the input to a linear transfer function $H(s)$ and applying the Laplace transform to (3.19), the prediction error dynamics are rewritten as

$$\tilde{x}(s) = H(s) \left(1 - C(s) - \Lambda C(s) \hat{\Lambda}^{-1} \right) \hat{\Omega}^\top \Phi(x), \quad (3.21)$$

where

$$H(s) = (sI - A_l(v, h))^{-1} B(v, h). \quad (3.22)$$

Applying the \mathcal{L}_∞ input–output bound established in Section 2.3.1, the peak prediction error is bounded as

$$\|\tilde{x}(s)\|_{\mathcal{L}_\infty} \leq \left\| H(s) \left(1 - C(s) - \Lambda C(s) \hat{\Lambda}^{-1} \right) \right\|_{\mathcal{L}_1} \left\| \hat{\Omega}^\top \Phi(x) \right\|_{\mathcal{L}_\infty}. \quad (3.23)$$

To obtain a worst-case bound independent of the estimated quantities, the estimated control effectiveness matrix $\hat{\Lambda}$ is replaced by its worst-case bound $\eta \|\Lambda\|_\infty$, and the estimated uncertainty $\hat{\Omega}^\top \Phi(x)$ is replaced by its worst-case bound $\|\Omega^\top \Phi(x)\|_\infty$. Requiring the resulting expression to be strictly less than unity, in accordance with the small-gain theorem [11], yields the \mathcal{L}_1 small-gain condition

$$\|H(s) \left(1 - C(s) - \|\Lambda\|_\infty C(s) (\eta \|\Lambda\|_\infty)^{-1} \right)\|_{\mathcal{L}_1} \|\Omega^\top \Phi(x)\|_\infty < 1, \quad (3.24)$$

where η is a scaling factor applied to the input uncertainty bound Λ_{\max} to suppress the \mathcal{L}_1 norm further. Consequently, selecting $\eta > 1$ effectively relaxes the projection constraint, allowing the estimate $\hat{\Lambda}$ to exceed the nominal bound and enter regions that may not be physically realistic. However, by selecting a sufficiently large projection tolerance ϵ in (2.50), the projection operator engages earlier as the estimate update approaches its constraint boundary, resulting in a more gradual attenuation of the update. This mitigates potential adverse effects introduced by the relaxed constraint. As a result, the dynamics of the predictor is able to be less aggressive while still maintaining a small enough \mathcal{L}_1 norm such that (3.27) holds.

To ensure (3.24) holds for any $\Phi(x) \in \mathcal{X}$, where \mathcal{X} denotes the feasible state space defined by the imposed state constraints, the regressor vector is normalized to satisfy

$$\|\Omega^\top \Phi(x)\|_\infty \leq \Omega_{\max}, \quad \|\Phi(x)\|_\infty \leq 1, \quad (3.25)$$

using

$$\Phi(x) = \frac{2(x - x_{\min})}{x_{\max} - x_{\min}} - 1, \quad (3.26)$$

where x_{\min} and x_{\max} are hard constraints which causes the simulation to stop when violated. The constraints for the integral states are obtained empirically using Monte Carlo simulation, which is further discussed in Section 3.6.3.

Thus, (3.24) is simplified to

$$\|H(s) (1 - C(s) - 1/\eta C(s))\|_{\mathcal{L}_1} \Omega_{\max} < 1. \quad (3.27)$$

The maximum uncertainty Ω_{\max} and Λ_{\max} is experimentally derived by simulating multiple uncertainty combinations and selecting the largest

$$\Omega_{\max} = \|B^\dagger A_\Delta\|_\infty, \quad \Lambda_{\max} = \|I + B^\dagger B_\Delta\|_\infty, \quad (3.28)$$

out of these simulations.

The predictor dynamics and η are jointly tuned to satisfy (3.27): the predictor poles are shifted further along the negative real axis relative to those of the reference model $A_m(\cdot)$, and η is evaluated across three candidate values, with the predictor poles retuned for each, until a low-pass filter exists such that (3.27) holds. Excessive model mismatch amplification is assessed by evaluating closed-loop stability in simulation, ensuring that the selected predictor dynamics do not introduce instability. This is evaluated for the pitch and yaw channels, while the roll channel uncertainty is sufficiently small such that the predictor dynamics match the reference model while still satisfying (3.27). The effect of insufficient η on closed-loop stability and the sensitivity of the \mathcal{L}_1 norm to this tuning procedure are illustrated in Section 4.1 and discussed further in Section 5.1.

3.5 Implemented \mathcal{L}_1 CMRAC Architecture

The implemented \mathcal{L}_1 composite MRAC architecture is shown in Figure 3.6 and is applied independently to the roll, pitch, and yaw channels. The architecture consists of two interconnected adaptive loops with distinct and complementary objectives. The inner loop, comprising the predictor model and the uncertainty adaptive law, is designed primarily for robustness: by employing fast adaptation rates, it rapidly estimates and attenuates the effect of uncertainties, ensuring the prediction error remains bounded via the \mathcal{L}_1 small-gain condition (2.48). The outer loop, comprising the reference model and the composite adaptive laws for the feedback and feedforward gains, is designed primarily for tracking performance: by employing slower adaptation rates, it keeps the adaptive gains close to their nominal values while penalizing reference tracking error, ensuring closed-loop stability and accurate reference tracking. The blue and red lines highlights the tracking and prediction error, respectively, fed into the respective adaptive parameter.

Thus, by extending the \mathcal{L}_1 adaptive controller described in Section 2.4 with an outer-loop, the implemented controller is able to handle tracking and robustness independently. This extension decouples robustness from tracking enabling fast uncertainty estimation while still maintaining accurate tracking performance similar to the baseline controller described in Section 3.1.3 under nominal conditions.

To verify the implementation and isolate the intrinsic behavior of the proposed control architecture, an additional evaluation is carried out using a simplified constant-velocity model without nonlinearities or cross-coupling. The results are presented in Appendix A.1.

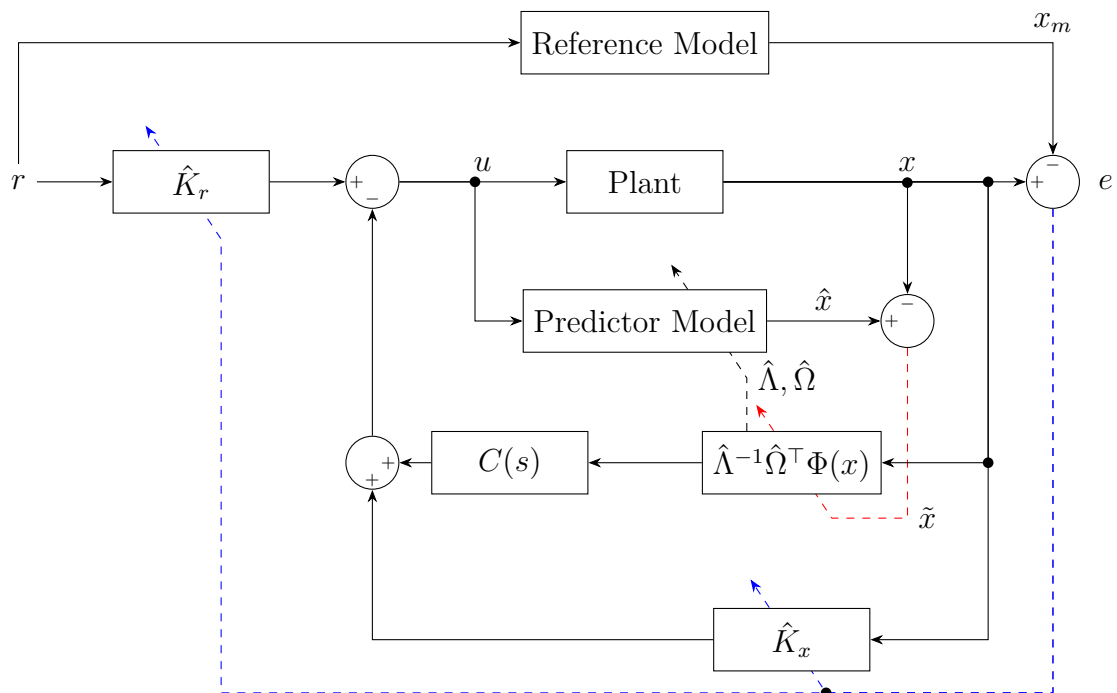


Figure 3.6: Block diagram of the \mathcal{L}_1 composite MRAC architecture. The plant dynamics in (2.6) are controlled to follow the reference model in (2.7). The control input (3.9) includes a strictly proper low-pass filter $C(s)$. The predictor (3.10) generates the estimate \hat{x} , forming the prediction error \tilde{x} , while the tracking error e is computed relative to the reference model. Both errors drive the composite adaptive laws, and the predictor dynamics are selected to satisfy the small-gain condition in (2.48), ensuring bounded-input bounded-output stability.

3.5.1 Tuning Parameters

The tuning parameters used for each channel, including the adaptation gains and Lyapunov weighting matrices, are summarized in Tables 3.2 and 3.3. These parameters determine the adaptation rate and influence the robustness and transient performance of the controller, and are obtained through a combination of empirical tuning and engineering intuition.

Due to the geometric symmetry of the vehicle, the pitch and yaw channels are tuned identically, reflecting the equivalent dynamic characteristics of these axes. The roll channel is tuned independently to account for its distinct dynamics.

To account for cross-coupling and nonlinearities across all three control channels, the regressor vector $\Phi(x)$ includes the full state vector

$$\Phi(x) = \begin{bmatrix} x_\varphi \\ x_\vartheta \\ x_\psi \end{bmatrix}, \quad (3.29)$$

where x_φ , x_ϑ , and x_ψ denote the roll, pitch, and yaw states defined in (3.5), respectively.

Several tuning choices warrant particular attention. The Lyapunov weighting matrix Q_m is selected to be sufficiently large for all control channels, ensuring that

the Lyapunov function (3.15) for the outer-loop is strictly decreasing given that the inner-loop remains stable, as required by the composite stability argument. The prediction error penalty Q_l is kept small relative to Q_m , reflecting the intentionally fast predictor dynamics. A large Q_l would over-penalize the prediction error, forcing the true states to closely follow the predicted states. Given the fast predictor dynamics, this would introduce oscillatory behavior in the adaptive update which may lead to instability. The adaptation rate Γ_Λ is selected to be significantly larger than the remaining adaptation rates, as seen in Tables 3.2 and 3.3, to prevent the estimate $\hat{\Lambda}$ from approaching zero, which would cause numerical issues due to the appearance of $\hat{\Lambda}^{-1}$ in the control law (3.9).

Tables 3.2 and 3.3 further presents the final values for the scaling factor η and the projection tolerance ϵ for each control channel. The importance of selecting $\eta > 1$ and its effect on closed-loop stability are illustrated in Section 4.1.

Table 3.2: Tuning parameters and their value for roll channel

Notation	Value
Γ_Λ	1×10^{10}
Γ_Ω	1×10^{-4}
Γ_r	5×10^{-8}
Γ_x	1×10^{-8}
Q_l	1×10^{-1}
Q_m	1×10^5
η	1
η_Λ	1.2
ϵ	0.1
ϵ_Λ	10

Table 3.3: Tuning parameters and their value for pitch and yaw channel

Notation	Value
Γ_Λ	1×10^{10}
Γ_Ω	5×10^{-1}
Γ_r	5×10^{-13}
Γ_x	1×10^{-12}
Q_l	5×10^{-3}
Q_m	2×10^5
η	1
η_Λ	10
ϵ	0.1
ϵ_Λ	10

3.6 Evaluation

To answer the research questions presented in Section 1.4, the proposed CMRAC controller is systematically compared against two baseline controllers: the nominal LQ controller described in Section 3.1.3 and a previously implemented direct MRAC [1]. This three-way comparison is designed to isolate the contribution of the composite adaptation law relative to both a non-adaptive and a purely error-driven adaptive benchmark.

Performance is assessed based on trajectory tracking performance across multiple reference paths, which is quantified through the root-mean-square error (RMSE) and maximum absolute error (MAE) computed for the acceleration tracking in pitch and yaw, while roll is assessed via roll angle tracking error. These metrics collectively capture both the average and worst-case deviations from the commanded accelerations and angle, providing a comprehensive evaluation of controller performance.

3.6.1 Paths

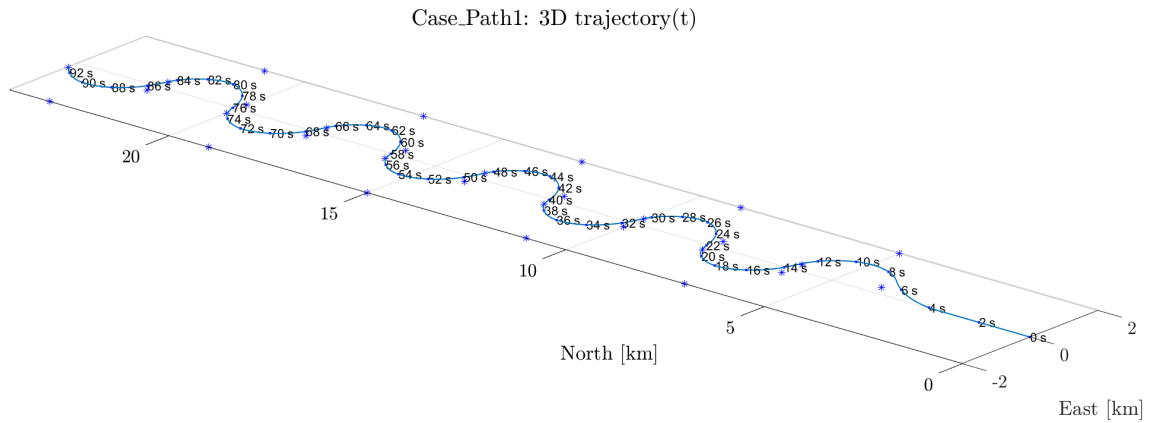
As mentioned in Section 3.1.2 and 3.1.3, the linearized models in (3.3) are derived under the assumption of decoupled channel dynamics. When multiple channels are simultaneously excited, this assumption may no longer hold, resulting in modeling errors due to non-linear cross-coupling effects. Evaluating the controllers under such conditions is therefore of particular interest, as it provides insight into their ability to mitigate model uncertainties and non-linear cross-coupling effects.

The first path is designed to primarily evaluate the controller's ability to track yaw commands. Figure 3.7a illustrates the corresponding way-points together with a representative trajectory generated by the LQ controller using nominal model parameters. The way-points form four sharp S-turn in the xy -plane. The corresponding acceleration in each control channel and Mach number is shown in Figure 3.7b and 3.7c, respectively.


The second path introduces more rapid changes in altitude, evaluating the controller ability to track pitch commands. The associated way-points and a representative nominal trajectory are shown in Figure 3.8a, where the points form four sharp S-turn in the xz -plane. The resulting accelerations and Mach profile for this scenario is presented in Figure 3.8b and 3.8c, respectively, where the same performance metrics are applied.

Lastly, the third path is designed to evaluate controller performance during combined multi-axis motion, requiring simultaneous tracking in multiple channels by following a path of three spiral turns, evaluating the controller in the event where non-linear cross-couplings are strong. The corresponding way-points and nominal trajectory are illustrated in Figure 3.9a, while the acceleration tracking results and corresponding Mach profile are shown in Figure 3.9b and 3.9c, respectively.

Together, these scenarios provide a comprehensive evaluation of the controller performance across different operating conditions, enabling a systematic comparison of tracking accuracy, robustness, and transient response.




(a) Way-points and nominal trajectory for sharp S-turns in xy -plane.




(b) Acceleration profile

(c) Mach profile of trajectory

Figure 3.7: Baseline evaluation for Path 1.




(a) Way-points and nominal trajectory for sharp S-turns in xy -plane.




(b) Acceleration profile

(c) Mach profile of trajectory

Figure 3.8: Baseline evaluation for Path 2.



(a) Way-points and nominal trajectory for sharp S-turns in xy -plane.



(b) Acceleration profile

(c) Mach profile of trajectory

Figure 3.9: Baseline evaluation for Path 3.

3.6.2 Predictor Evaluation

The controller is further evaluated by designing three predictors corresponding to different uncertainty bounds. The objective is to assess the sensitivity of the predictor design to the assumed uncertainty limits and to verify satisfaction of the \mathcal{L}_1 small-gain condition. The \mathcal{L}_1 norm is evaluated at the four operating points illustrated in Figure 3.3.

The nominal predictor is designed using the experimentally obtained bound without modification, such that the \mathcal{L}_1 norm at each operating point lies on the boundary of the \mathcal{L}_1 small-gain condition. The resulting norms are shown in Figure 3.10. While all operating points satisfy the \mathcal{L}_1 small-gain condition in (2.48), interpolation between these points leads to violations in certain regions, particularly in the yaw channel, as shown in Figures 3.10e and 3.10f.

This behavior highlights an additional limitation of the model derived in Section 3.1.2, where interpolation between linearized operating points may, in addition to cross-coupling, introduce non-linear effects. Therefore, evaluating predictors tuned both below and above the boundary provides insight into the robustness of the design with respect to interpolation-induced nonlinearities.

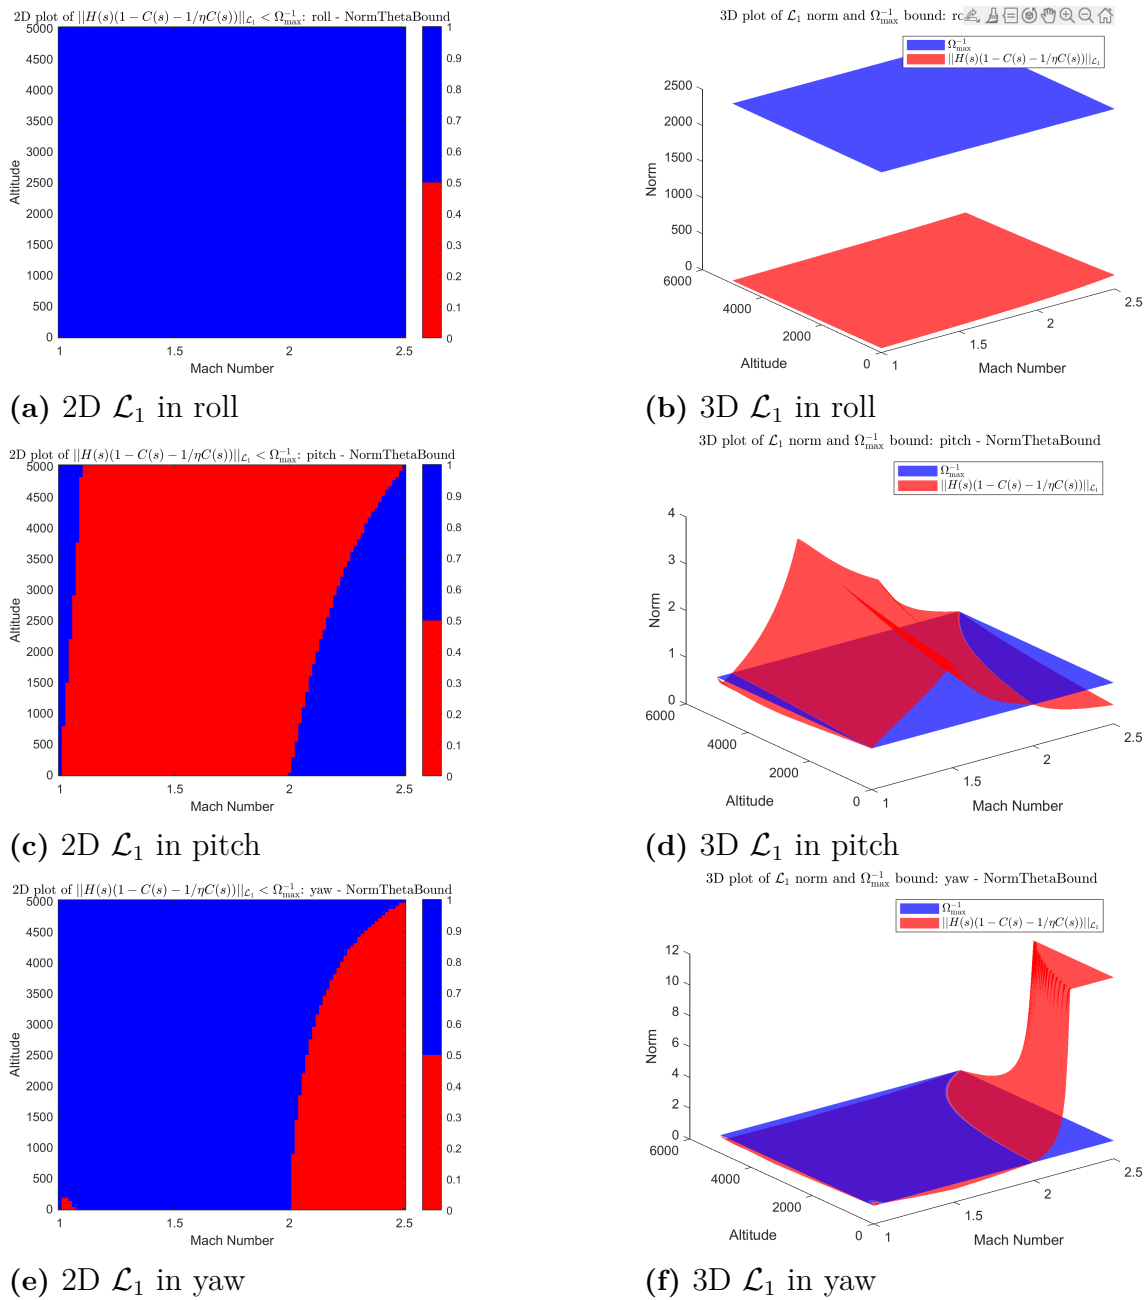


Figure 3.10: \mathcal{L}_1 norms produced with predictor tuned for nominal Ω_{\max} .

3.6.3 Monte Carlo

To ensure statistical robustness and to assess controller behaviour under a broad range of uncertainty conditions, a Monte Carlo (MC) simulation study [19] is conducted for each controller on each path. Each MC campaign randomly samples from the defined uncertainty set, and performance is further quantified by the success rate, defined as the proportion of trials in which the vehicle successfully clears the path. The number of MC simulations is derived using Chebyshev's inequality

3. Methods

[20]

$$N \geq \frac{\hat{p}(1 - \hat{p})}{\delta\epsilon^2}, \quad (3.30)$$

where N is the number of simulations required, \hat{p} is the estimated probability of success, δ is the significance level such that $1 - \delta$ is the confidence level, and ϵ is the desired estimation accuracy, i.e., the maximum allowable deviation between the estimated and true success probability. By taking into account computational complexity, the confidence level δ is set to 5%, and with no prior belief of success rate \hat{p} is set to 50% with a maximum allowed deviation of 10%. Thus, $N = 500$.

The normalization applied to the regressor vector $\Phi(x)$ in (3.26) is performed using hard bounds for all states defined in (3.5), with the exception of the integral states. For these states, the normalization limits are instead determined empirically based on 500 MC simulations conducted across all paths using the LQ controller. These simulations also provide an initial prior estimate of the success rate, which are summarized in Table 3.4, along with the recalculated number of Monte Carlo simulations required, rounded up to the nearest multiple of 10. The plotted MC traces for the three paths are illustrated in Figure 3.11–3.13. To ensure a fair comparison, all MC simulations were performed using identical random samples of the parameters for each controller.

Table 3.4: A priori success rate estimated from initial Monte Carlo simulations.

Path	Estimated Success Rate \hat{p}	Monte Carlo Simulations N
Path 1	$\approx 74\%$	390
Path 2	$\approx 86\%$	250
Path 3	$\approx 88\%$	220

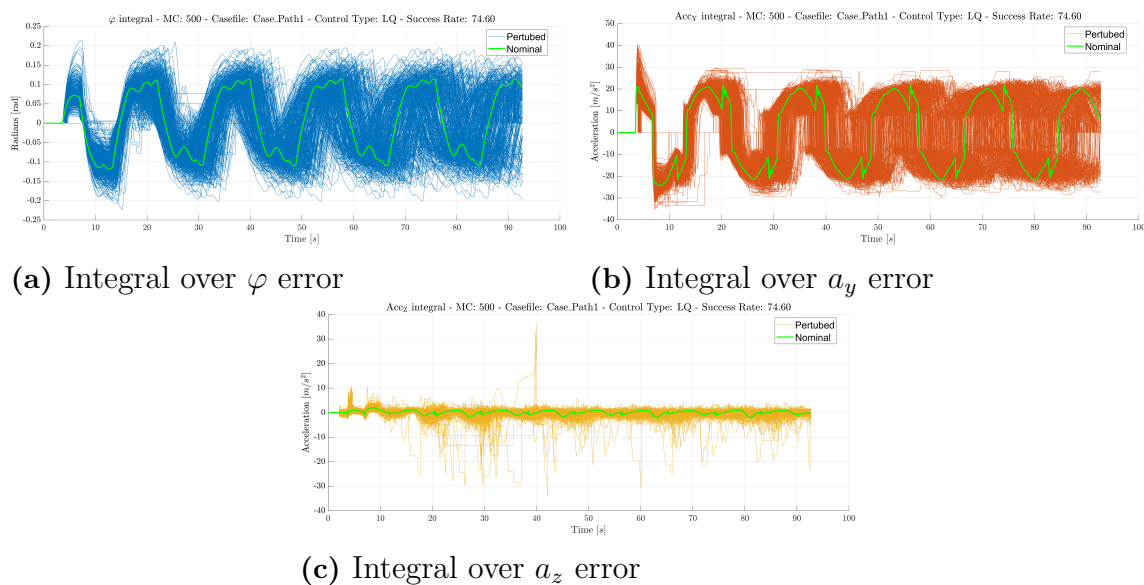


Figure 3.11: Monte Carlo traces of the integral states for Path 1 obtained from 500 simulations using the LQ controller.

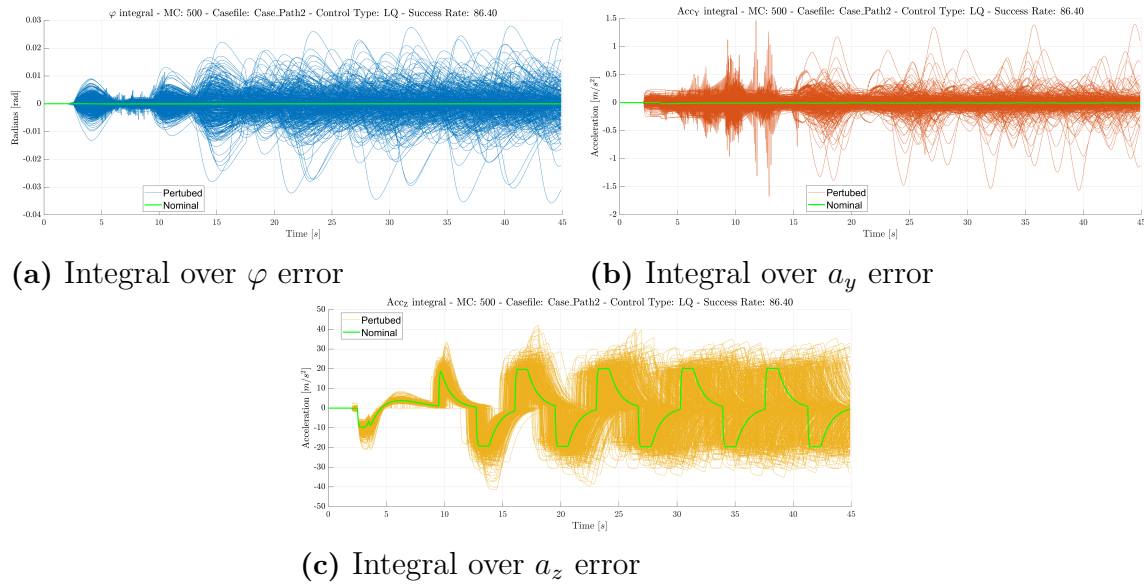


Figure 3.12: Monte Carlo traces of the integral states for Path 2 obtained from 500 simulations using the LQ controller.

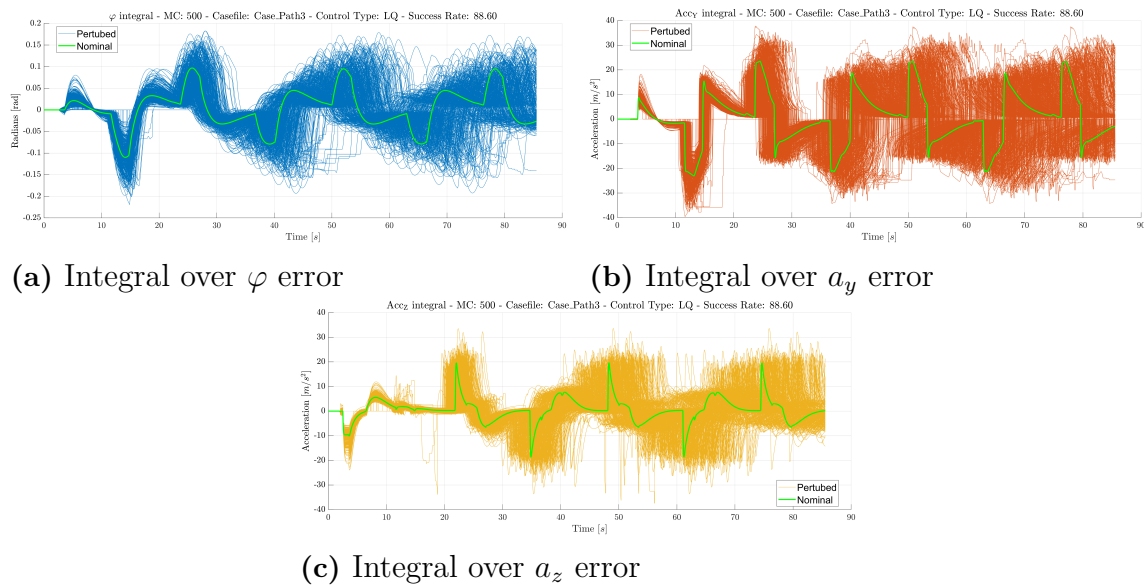


Figure 3.13: Monte Carlo traces of the integral states for Path 3 obtained from 500 simulations using the LQ controller.

3.6.4 Sensitivity Analysis

To assess controller sensitivity to external disturbances, Path 3 is reevaluated for nominal parameter values under the influence of an exogenous input. The disturbance, representing a wind gust, is applied perpendicular to the velocity vector at a corner of the trajectory where both the pitch and yaw channels are simultaneously excited. The disturbance direction relative to the velocity vector is illustrated in

3. Methods

Figure 3.14, where its magnitude is set to 1.5×32.7 m/s, corresponding to 50% over hurricane-level wind speeds [21]. This evaluation is performed on the LQ, MRAC, and the best performing \mathcal{L}_1 CMRAC.

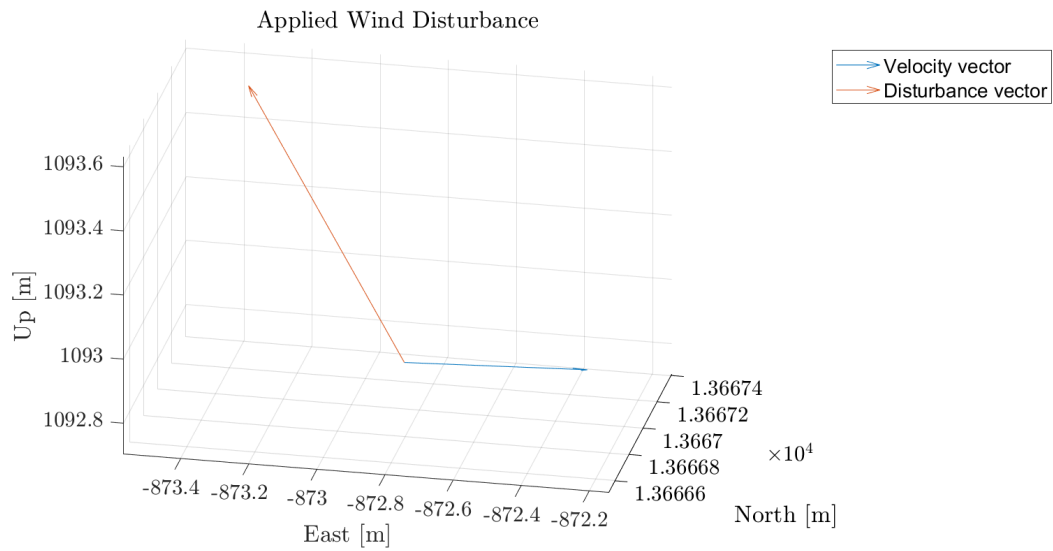


Figure 3.14: Applied disturbance perpendicular to velocity vector of vehicle.


4

Results


This chapter is introduced by motivating the importance of the scaling factor η . This is followed by presenting the \mathcal{L}_1 norm in (3.27) for three different predictor dynamics. Lastly, results from each controller is given.

4.1 Scaling Factor Tuning


The values for η and ϵ given in Tables 3.2 and 3.3 are motivated by comparing three cases of adaptive parameters. In each case, the poles of the closed-loop predictor, initialized such that (3.27) holds for $\eta = 1$, is shifted closer to the origin along the negative real axis while the scaling factor is increased. Thus, for each case, the \mathcal{L}_1 small gain holds for conservative bound, while bounded prediction error may not. The results of these three cases are presented in Figure 4.1 where the controller is evaluated on path 1 with nominal model parameters.



(a) $\eta = 1$



(b) $\eta = 5$



(c) $\eta = 10$

Figure 4.1: Adaptive parameters for different scaling factors η on Path 1.

4.2 \mathcal{L}_1 Norm

As mentioned in Section 3.6.2, three predictor configurations are designed, including a nominal case. The first configuration is obtained by conservatively increasing the experimentally identified uncertainty bound Ω_{\max} by 20%. A larger uncertainty bound requires faster predictor dynamics in order to satisfy the \mathcal{L}_1 small-gain condition in (3.27), yielding a more aggressive predictor. Consequently, the resulting \mathcal{L}_1 norms are reduced.

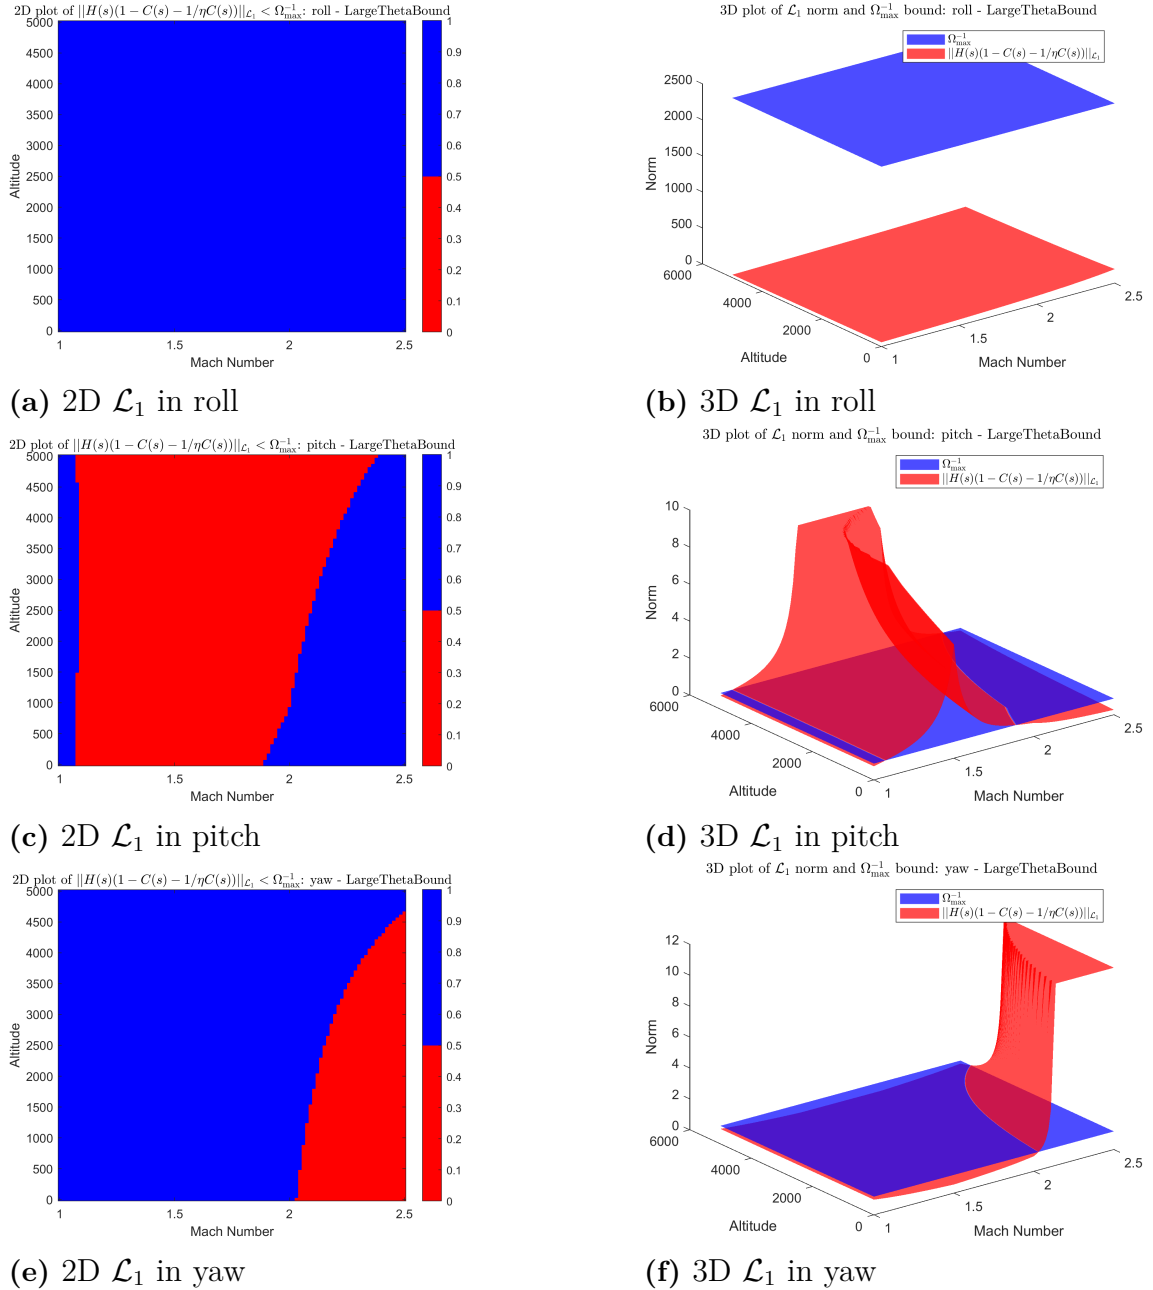


Figure 4.2: \mathcal{L}_1 norms produced with predictor tuned for $1.2 \times \Omega_{\max}$.

The corresponding 2D and 3D \mathcal{L}_1 norms for each rotational channel relative to the

inverse of their respective bounds are illustrated in Figure 4.2, where the \mathcal{L}_1 norm for all operating points in Figure 3.3 satisfy (3.27).

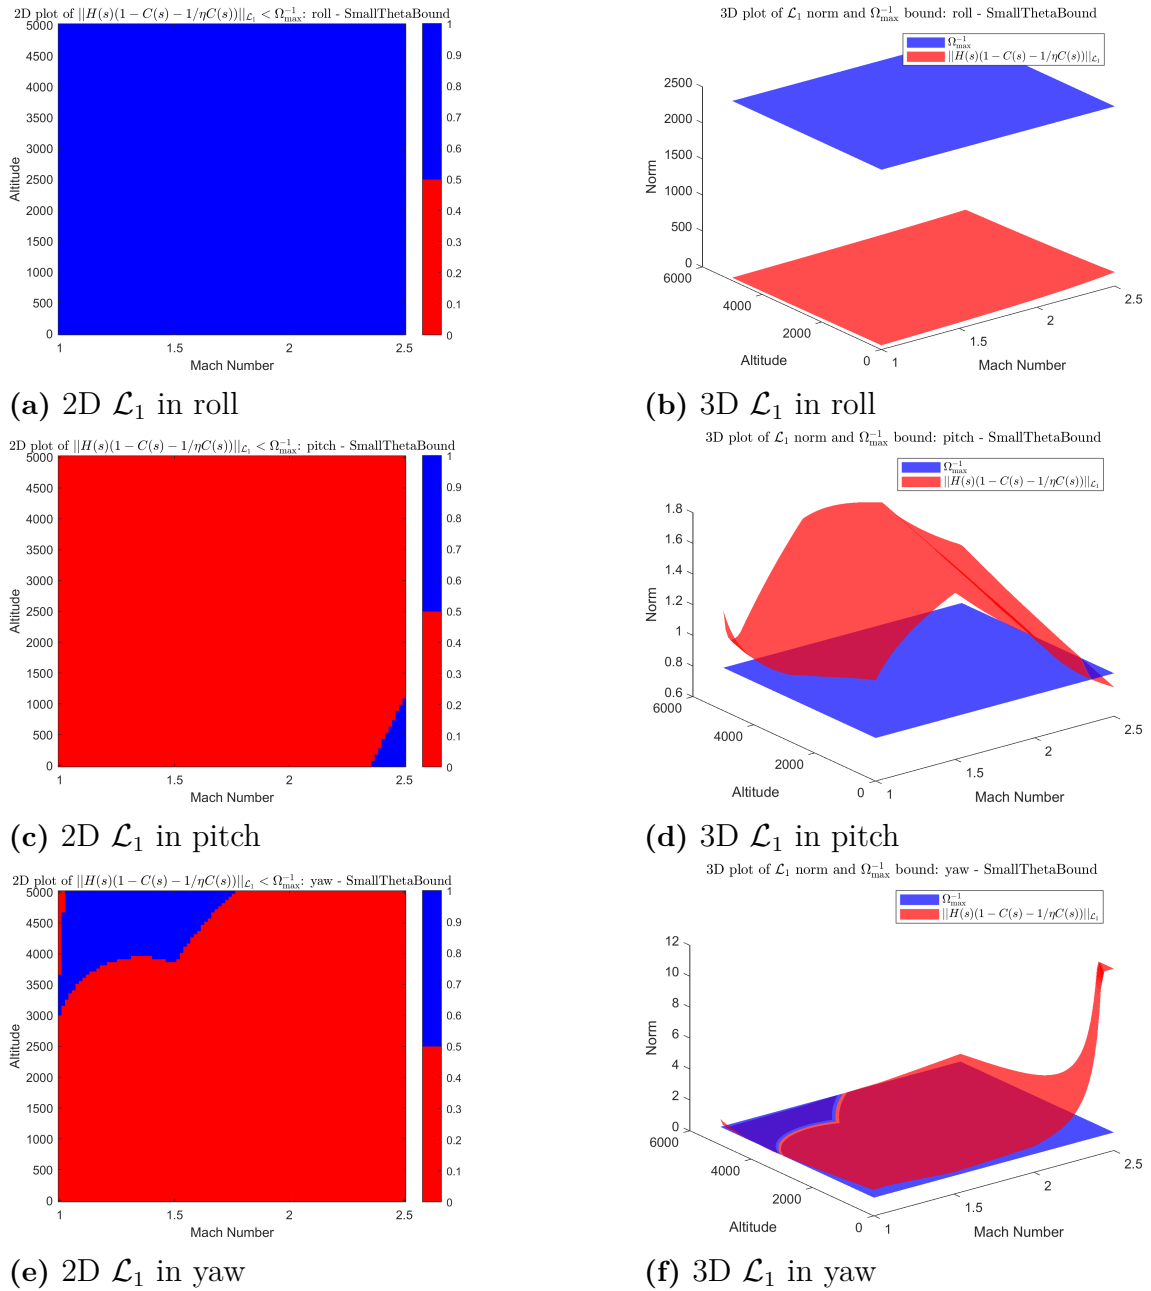


Figure 4.3: \mathcal{L}_1 norms produced with predictor tuned for $0.7 \times \Omega_{\max}$.

A conservative predictor is designed by reducing the experimental bound by 30%. In this case, slower predictor dynamics are sufficient to satisfy the small-gain condition, resulting in larger \mathcal{L}_1 norms. The corresponding results are illustrated in Figure 4.3 were for both the pitch and yaw channel, all points within the polytope shown in Figure 3.3 violate (3.27).

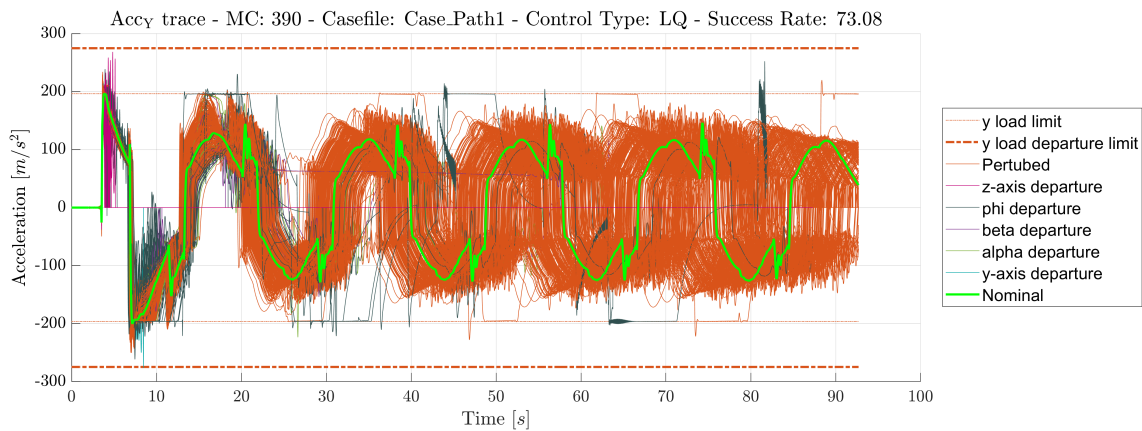
4.3 Controller Performance

4.3.1 Path 1


In this section, the results obtained for Path 1 are presented. The performance of the LQ, MRAC, and CMRAC controllers are evaluated using Monte Carlo simulations.

Linear Quadratic Regulator

The performance of the LQ controller for Path 1 is illustrated in Figures 4.4 and 4.5. The figures show the acceleration and roll angle responses relative to the imposed constraints and reference trajectories, respectively.



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.4: Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the LQ controller.

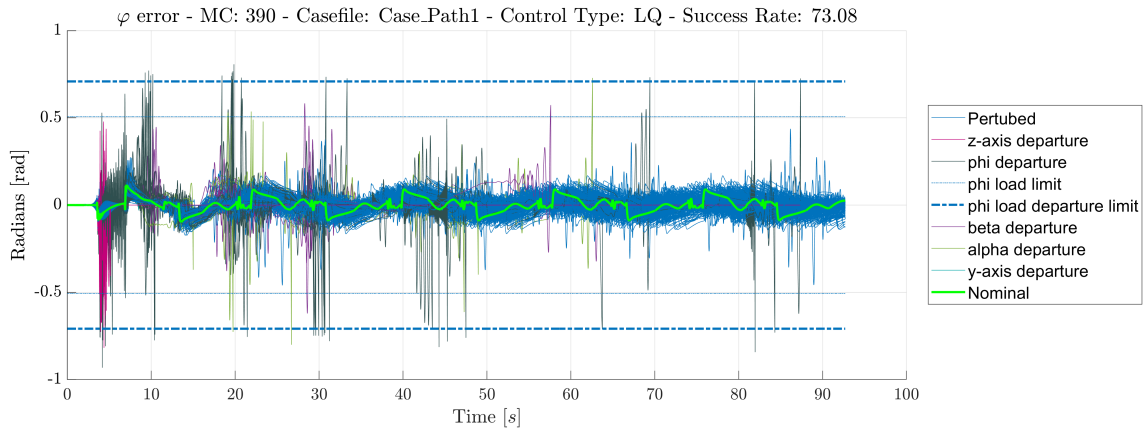
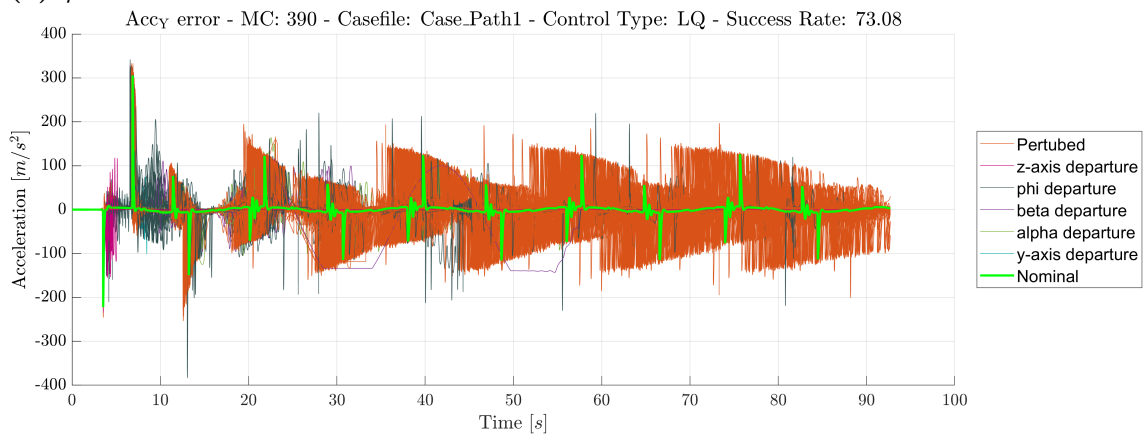


(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference


Figure 4.5: Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the LQ controller.

Model Reference Adaptive Controller

The performance of the MRAC for Path 1 is illustrated in Figures 4.6 and 4.7. The figures show the acceleration and roll angle responses relative to the imposed constraints and reference trajectories, respectively.



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.6: Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the MRAC controller.





(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference


Figure 4.7: Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the MRAC controller.

\mathcal{L}_1 Composite Model Reference Adaptive Controller

The performance of the \mathcal{L}_1 CMRAC for Path 1 is evaluated for three different uncertainty bounds. These correspond to the nominal, large, and small uncertainty configurations derived in Section 3.6.2. The resulting Monte Carlo simulations are presented in Figures 4.8–4.13.



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.8: Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.






(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference

Figure 4.9: Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.

4. Results



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.10: Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.




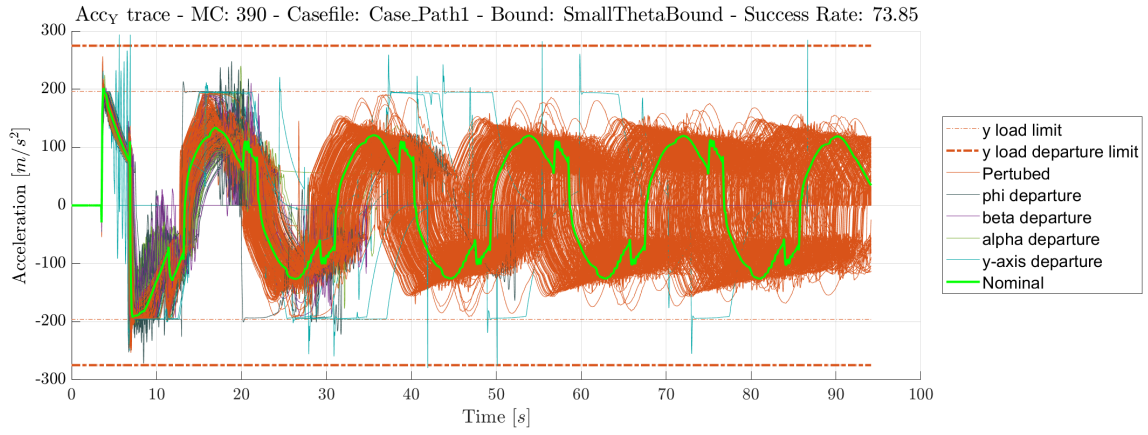

(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference

Figure 4.11: Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.

4. Results



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.12: Monte Carlo traces relative to constraints for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.



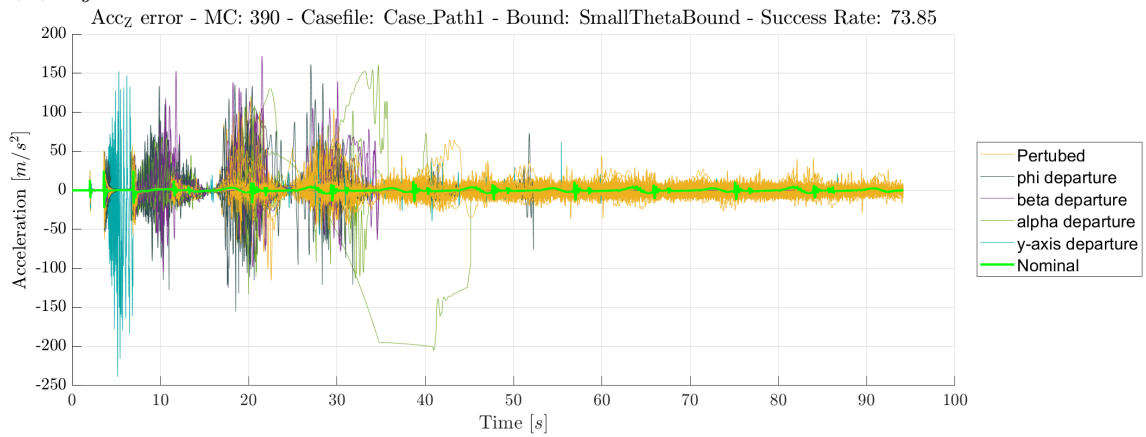
(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference

Figure 4.13: Monte Carlo tracking errors for Path 1 obtained from 390 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.

Performance Metric

The average MAE and RMSE for all successful MC simulations for Path 1 are presented in Table 4.1–4.3 for the roll, pitch, and yaw channel respectively. The tables also summarize the success rate for each controller. The best and worst cases are highlighted with green and red colors, respectively.

Table 4.1: Performance results of each controller for roll channel in Path 1. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.

Controller	MAE	RMSE	Success Rate
LQR	1.293×10^{-1}	2.948×10^{-2}	73.08%
MRAC	1.584×10^{-1}	4.271×10^{-2}	76.77%
NT CMRAC	1.266×10^{-1}	3.146×10^{-2}	73.08%
LT CMRAC	1.270×10^{-1}	3.142×10^{-2}	74.10%
ST CMRAC	1.256×10^{-1}	3.131×10^{-2}	73.85%

Table 4.2: Performance results of each controller for pitch channel in Path 1. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.

Controller	MAE	RMSE	Success Rate
LQR	40.43	2.786	73.08%
MRAC	32.26	4.912	76.77%
NT CMRAC	33.76	2.409	73.08%
LT CMRAC	33.87	2.428	74.10%
ST CMRAC	33.49	2.386	73.85%

Table 4.3: Performance results of each controller for yaw channel in Path 1. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.


Controller	MAE	RMSE	Success Rate
LQR	302.6	15.37	73.08%
MRAC	298.3	16.47	76.77%
NT CMRAC	303.6	15.48	73.08%
LT CMRAC	303.5	15.47	74.10%
ST CMRAC	303.7	15.47	73.85%

4.3.2 Path 2


This section presents the results obtained for Path 2. The performance of the LQ, MRAC, and \mathcal{L}_1 CMRAC controllers is evaluated using Monte Carlo simulations. The responses are analyzed in terms of constraint satisfaction and tracking performance.

Linear Quadratic Regulator

The performance of the LQ controller for Path 2 is illustrated in Figures 4.14 and 4.15. The figures show the roll angle and acceleration responses relative to the imposed constraints and the corresponding tracking errors with respect to the reference trajectory.




(a) a_y relative to constraints




(b) a_z relative to constraints

Figure 4.14: Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the LQ controller.


4. Results



(a) φ relative to reference



(b) a_y relative to reference




(c) a_z relative to reference


Figure 4.15: Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the LQ controller.

Model Reference Adaptive Controller

The performance of the MRAC for Path 2 is presented in Figures 4.16 and 4.17. The results illustrate the constraint handling capabilities and tracking performance under the considered uncertainty realizations.




(a) a_y relative to constraints




(b) a_z relative to constraints

Figure 4.16: Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the MRAC controller.

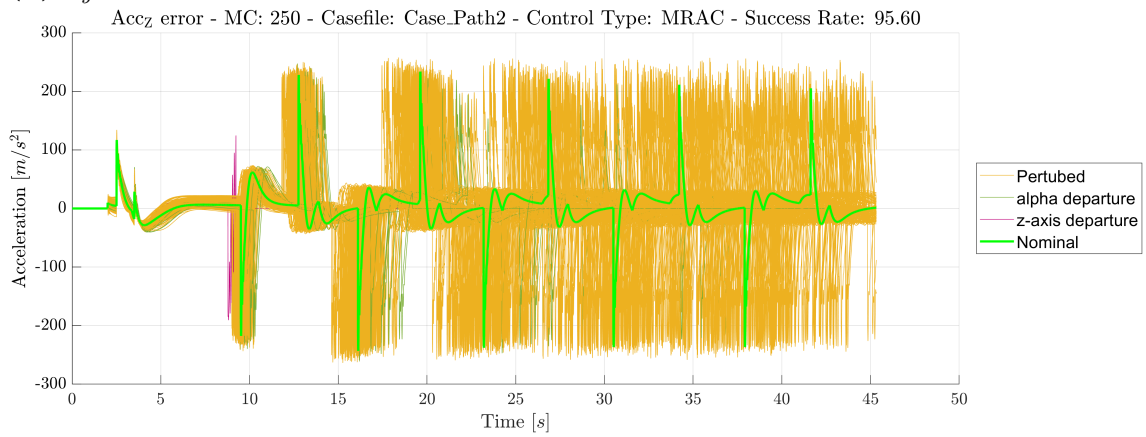
4. Results



(a) φ relative to reference



(b) a_y relative to reference




(c) a_z relative to reference


Figure 4.17: Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the MRAC controller.

\mathcal{L}_1 Composite Model Reference Adaptive Controller

The performance of the \mathcal{L}_1 CMRAC for Path 2 is evaluated for three different uncertainty bounds. These correspond to the nominal, large, and reduced uncertainty configurations derived in Section 3.6.2. The resulting Monte Carlo simulations are presented in Figures 4.18–4.23.




(a) a_y relative to constraints




(b) a_z relative to constraints

Figure 4.18: Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.


4. Results



(a) φ relative to reference



(b) a_y relative to reference



(c) a_z relative to reference

Figure 4.19: Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.



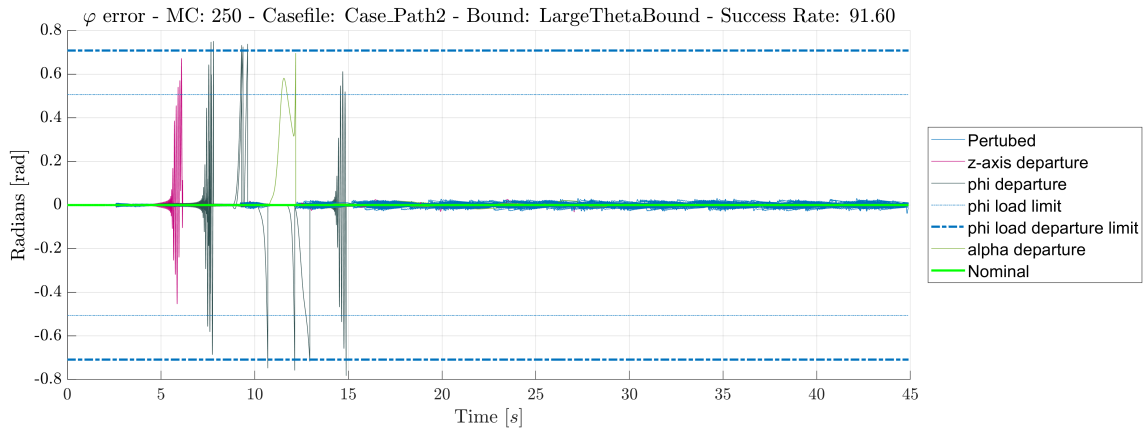

(a) a_y relative to constraints(b) a_z relative to constraints

Figure 4.20: Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.


4. Results



(a) φ relative to reference



(b) a_y relative to reference



(c) a_z relative to reference

Figure 4.21: Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.





(a) a_y relative to constraints(b) a_z relative to constraints

Figure 4.22: Monte Carlo traces relative to constraints for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.


4. Results



(a) φ relative to reference



(b) a_y relative to reference



(c) a_z relative to reference

Figure 4.23: Monte Carlo tracking errors for Path 2 obtained from 250 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.

Performance Metric

For Path 2, the average MAE and RMSE computed over all successful MC simulations are presented in Tables 4.4–4.6 for the roll, pitch, and yaw channels, respectively. The tables also report the success rate for each controller. The best and worst cases are highlighted with green and red colors, respectively.

Table 4.4: Performance results of each controller for roll channel in Path 2. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.

Controller	MAE	RMSE	Success Rate
LQ	1.155×10^{-2}	1.893×10^{-3}	86.40%
MRAC	1.095×10^{-2}	2.625×10^{-3}	95.60%
NT CMRAC	9.788×10^{-3}	1.678×10^{-3}	94.00%
LT CMRAC	9.813×10^{-3}	1.691×10^{-3}	91.60%
ST CMRAC	9.807×10^{-3}	1.679×10^{-3}	94.40%

Table 4.5: Performance results of each controller for pitch channel in Path 2. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.

Controller	MAE	RMSE	Success Rate
LQ	233.0	15.44	86.40%
MRAC	242.9	19.72	95.60%
NT CMRAC	239.4	14.60	94.00%
LT CMRAC	239.1	14.64	91.60%
ST CMRAC	239.4	14.60	94.40%

Table 4.6: Performance results of each controller for yaw channel in Path 2. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.


Controller	MAE	RMSE	Success Rate
LQ	2.280	1.335×10^{-1}	86.40%
MRAC	8.906	1.162	95.60%
NT CMRAC	1.909	1.078×10^{-1}	94.00%
LT CMRAC	1.901	1.078×10^{-1}	91.60%
ST CMRAC	1.908	1.081×10^{-1}	94.40%

4.3.3 Path 3


This section presents the results obtained for Path 3. The LQ, MRAC, and \mathcal{L}_1 CMRAC controllers are evaluated using Monte Carlo simulations to assess tracking performance and constraint satisfaction under varying uncertainty realizations.

Linear Quadratic Regulator

The performance of the LQ controller for Path 3 is shown in Figures 4.24 and 4.25. The figures illustrate the roll angle and acceleration responses relative to constraints, as well as the corresponding tracking errors with respect to the reference trajectory.



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.24: Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the LQ controller.





(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference


Figure 4.25: Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the LQ controller.

Model Reference Adaptive Controller

The performance of the MRAC for Path 3 is presented in Figures 4.26 and 4.27. The results demonstrate the controller behavior in terms of constraint satisfaction and reference tracking under the considered uncertainty realizations.



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.26: Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the MRAC controller.

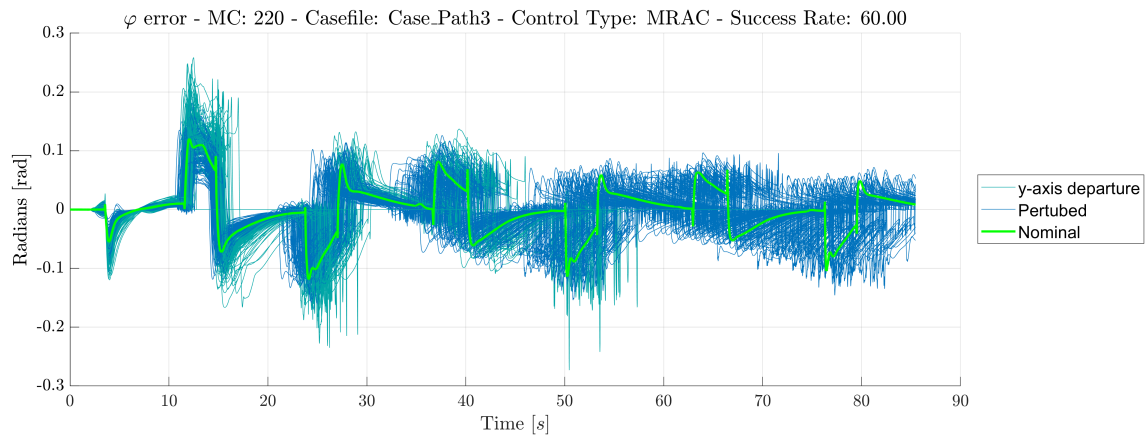



(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference


Figure 4.27: Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the MRAC controller.

\mathcal{L}_1 Composite Model Reference Adaptive Controller

The performance of the \mathcal{L}_1 CMRAC for Path 3 is evaluated for three different uncertainty bounds. These correspond to the nominal, large, and reduced uncertainty configurations derived in Section 3.6.2. The resulting Monte Carlo simulations are presented in Figures 4.28–4.33.



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.28: Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.






(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference

Figure 4.29: Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with nominal uncertainty bounds.

4. Results



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.30: Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.






(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference

Figure 4.31: Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with large uncertainty bounds.

4. Results



(a) a_y relative to constraints



(b) a_z relative to constraints

Figure 4.32: Monte Carlo traces relative to constraints for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.




(a) φ relative to reference(b) a_y relative to reference(c) a_z relative to reference

Figure 4.33: Monte Carlo tracking errors for Path 3 obtained from 220 simulations using the \mathcal{L}_1 CMRAC controller with reduced uncertainty bounds.

Performance Metric

For Path 3, the average MAE and RMSE over all successful MC simulations are summarized in Tables 4.7–4.9 for the roll, pitch, and yaw channels, respectively. The tables further include the success rate for each controller. The best and worst cases are highlighted with green and red colors, respectively.

Table 4.7: Performance results of each controller for roll channel in Path 3. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.

Controller	MAE	RMSE	Success Rate
LQ	1.187×10^{-1}	2.267×10^{-2}	86.82%
MRAC	1.118×10^{-1}	2.746×10^{-2}	60.00%
NT CMRAC	1.240×10^{-1}	2.527×10^{-2}	91.82%
LT CMRAC	1.249×10^{-1}	2.518×10^{-2}	91.82%
ST CMRAC	1.247×10^{-1}	2.520×10^{-2}	92.73%

Table 4.8: Performance results of each controller for pitch channel in Path 3. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.

Controller	MAE	RMSE	Success Rate
LQ	217.5	11.32	86.82%
MRAC	233.9	16.19	60.00%
NT CMRAC	226.6	11.56	91.82%
LT CMRAC	227.9	11.27	91.82%
ST CMRAC	227.1	11.45	92.73%

Table 4.9: Performance results of each controller for yaw channel in Path 3. NT, LT, and ST denote the predictor tuning for nominal, large, and small uncertainty bound, respectively.


Controller	MAE	RMSE	Success Rate
LQ	242.1	17.13	86.82%
MRAC	252.7	19.37	60.00%
NT CMRAC	246.3	17.57	91.82%
LT CMRAC	246.4	17.41	91.82%
ST CMRAC	247.0	17.43	92.73%

4.4 Controller Performance in the Presence of External Uncertainty


Since the \mathcal{L}_1 CMRAC tuned for $0.7 \times \Omega_{\max}$ achieved on average a higher success rate than the controllers equipped with nominal and conservative closed-loop predictor dynamics, it is selected for further evaluation under external disturbances and compared with the LQ and MRAC controllers.

The disturbance is applied at $t = 22$ s and remains active until $t = 25$ s. Figure 4.34 illustrates the evolution of the adaptive regressor weights $\hat{\Omega}$ for the MRAC and \mathcal{L}_1 CMRAC controllers during this interval, providing insight into how each controller adapts to the disturbance.

The resulting closed-loop response is shown in Figure 4.35, where the roll angle φ and the lateral and vertical accelerations (a_y and a_z) are presented for all three controllers. The corresponding nominal trajectories are included to enable a direct assessment of tracking disturbance rejection.




(a) Adaptive weights $\hat{\Theta}$ for MRAC




(b) Adaptive weights $\hat{\Omega}$ for \mathcal{L}_1 CMRAC

Figure 4.34: Adaptive parameter evolution under external disturbance for Path 3. The disturbance is active between $t = 22$ s and $t = 25$ s.


4. Results



(a) Roll angle φ



(b) Lateral acceleration a_y



(c) Vertical acceleration a_z

Figure 4.35: Closed-loop response under external disturbance for Path 3. The figure shows the roll angle and acceleration tracking for the LQ, MRAC, and \mathcal{L}_1 CMRAC controllers, together with the corresponding nominal trajectories. The disturbance is active between $t = 22$ s and $t = 25$ s.

5

Discussion

This chapter contains the the analysis regarding the results provided in Chapter 4.

5.1 Importance of Scaling Factor η for Input Uncertainty Bound Λ_{\max}

As shown in Section 4.1, the adaptive parameter behavior is strongly influenced by the aggressiveness of the closed-loop predictor dynamics. In Figure 4.1a, where $\eta = 1$ requires significantly faster predictor poles to satisfy the small-gain condition, the adaptive parameters begin to oscillate at approximately 2.2 s. In contrast, the adaptive parameters in Figure 4.1b, corresponding to $\eta = 5$, remain comparatively well-behaved until approximately 2.5 s, indicating improved robustness with less aggressive predictor dynamics. This trend is further illustrated in Figure 4.1c, which employs $\eta = 10$, where no oscillatory behavior is observed, demonstrating a substantial improvement in stability.

These results indicate that, although larger values of η relax the bound on the input uncertainty and allow $\hat{\Lambda}$ to enter regions that may be physically unrealistic, the controller maintains stable behavior in the adaptive parameters, particularly in $\hat{\Omega}$. This highlights that stability of the adaptive architecture can be preserved even under conservative or inflated uncertainty bounds.

The results also reveal adverse effects introduced by interpolation between operating points. When the predictor dynamics are overly aggressive, model mismatch is amplified, causing the prediction error \tilde{x} to grow unbounded. Consequently, a fundamental trade-off emerges: the predictor must be sufficiently fast to satisfy the \mathcal{L}_1 small-gain condition in (3.27), while not so aggressive that interpolation-induced modeling errors destabilize the predictor.

5.2 Validation of \mathcal{L}_1 Norms

In addition to the instability observed in Section 5.1, the \mathcal{L}_1 norm evaluation provides independent evidence that interpolation-induced mismatch amplification persists even under stable predictor configurations. Although the \mathcal{L}_1 small-gain condition in (3.27) is satisfied at the designated operating points shown in Figure 3.3, it is not guaranteed to hold throughout the entire polytope. In particular, violations occur in the regions between operating points due to interpolation of the linearized models, as interpolation introduces non-linear behavior that is not captured by the local

linear models, as observed in Figure 3.10 and discussed in Section 3.6.2. This effect is evident in Figures 4.2c–4.2d, corresponding to the pitch channel with aggressive predictor tuning.

A comparison between Figures 4.2 and 4.3 further quantifies the impact of predictor aggressiveness on the \mathcal{L}_1 norm elevation. As the predictor dynamics are made faster, the deviation between interpolated models is amplified, resulting in increased nonlinearities across the polytope. This is particularly pronounced in the pitch channel, where the maximum \mathcal{L}_1 norm exceeds $12 \times \Omega_{\max}$ for the aggressive predictor, as shown in Figure 4.2d, compared to approximately $1.6 \times \Omega_{\max}$ for the conservative predictor in Figure 4.3d. Together with the instability observed in Section 5.1, these results consistently support the conclusion that aggressive predictor dynamics are detrimental when gain scheduling is employed, and that a less aggressive predictor design is preferable to limit interpolation-induced mismatch amplification across the flight envelope.

Consequently, the elevated \mathcal{L}_1 norm between operating points directly corresponds to reduced robustness margins in those regions, as a larger norm implies a less restrictive bound on the effect of uncertainty on the prediction error dynamics.

5.3 Overall Performance Comparison

The results presented in Chapter 4 further support the observations made in Section 5.1, namely that increasingly fast closed-loop predictor dynamics tend to degrade overall controller performance. Table 5.1 summarizes success rates and the average success rate across all three paths.

Although the three \mathcal{L}_1 CMRAC configurations exhibit comparable performance, the controller tuned with less aggressive predictor dynamics achieves the highest average success rate. In contrast, the configuration based on a conservative uncertainty bound yields the lowest average performance among the \mathcal{L}_1 CMRAC controllers.

While the robust adaptive controllers overall gave the highest success rate, the performance-oriented MRAC exhibits the lowest overall success rate among all five controllers. This outcome is primarily driven by its poor performance in Path 3, where nonlinearities and cross-coupling effects are more pronounced. A more detailed, path-dependent analysis is provided in Section 5.3.1.

Table 5.1: Summarized success rates with controllers ranked from highest to lowest based on average success rate. ST, NT, and LT denote the predictor tuning for small, nominal, and large uncertainty bound, respectively.

Controller	Path 1	Path 2	Path 3	Average Success Rate
ST CMRAC	73.85%	94.40%	92.73%	86.99%
NT CMRAC	73.08%	94.00%	91.82%	86.30%
LT CMRAC	74.10%	91.60%	91.82%	85.84%
LQ	73.08%	86.40%	86.82%	82.10%
MRAC	76.77%	95.60%	60.00%	77.42%

5.3.1 Path Dependent Performance

Path 1

The performance metrics for Path 1, summarized in Table 4.1–4.3, indicate that all controllers achieve comparable tracking accuracy in terms of RMSE and MAE, with only minor variations. Among the evaluated controllers, the MRAC attains the highest success rate, whereas the LQ controller yields the lowest. However, this improvement in success rate comes at the expense of tracking performance, as the MRAC produces the highest RMSE across the control channels, while the LQ controller achieves the lowest RMSE, particularly in the roll and yaw channels.

Compared to the remaining paths, the overall success rate for Path 1 is noticeably lower across all controllers, indicating that this scenario is inherently more challenging. This increased difficulty can be attributed to two compounding factors.

First, the low flight velocity along Path 1, shown in Figure 3.7c, reduces the dynamic pressure acting on the control surfaces, diminishing their effectiveness and increasing the tendency to overshoot during the sharp corners of the trajectory. Each overshoot results in an unfavorable vehicle orientation at the entry of the subsequent turn, potentially requiring large corrective rotations approaching 180° about the z -axis to realign the vehicle with the subsequent path segment. These corrective maneuvers demand large lateral accelerations that induce significant sideslip angles, with the sideslip constraint being particularly critical as sudden and large lateral acceleration demands significantly increase the likelihood of exceeding allowable β limits, resulting in constraint violation.

Second, since the pitch channel is largely occupied with counteracting gravitational effects to maintain level flight, the available pitch authority for trajectory tracking is reduced, increasing cross-coupling between the pitch and yaw channels and further degrading the vehicle's ability to execute precise coordinated maneuvers. Together, these factors render Path 1 the most demanding of the evaluated scenarios, and the reduced success rates observed across all controllers reflect this inherent difficulty rather than a deficiency of any particular control design.

Path 2

The higher success rates observed in Path 2 compared to Path 1 can be attributed to the motion occurring primarily in the same plane as the gravitational force, i.e., the xz -plane, and the significantly increased velocity, as shown in Figure 3.8c. Consequently, the dominant excitation is introduced in the pitch channel, while the roll and yaw channel remains less excited, reducing cross-coupling effects.

The MC traces of the lateral and vertical acceleration for LQ, MRAC, \mathcal{L}_1 CMRAC equipped with nominal predictor, aggressive \mathcal{L}_1 CMRAC, and conservative \mathcal{L}_1 CMRAC are shown in Figure 4.14, 4.16, 4.18, 4.20, and 4.22, respectively. While all controllers achieve comparable success rates, a clear distinction is observed in the yaw channel, where both the LQ and MRAC controllers exhibit noticeable excitation in lateral acceleration. In particular, the MRAC results in a significantly increased RMSE, approximately $10\times$ higher than the other controllers. In contrast, all \mathcal{L}_1 CMRAC configurations effectively attenuate these effects and maintain a well-regulated

lateral acceleration response.

Further insight is obtained from the roll angle evolution shown in Figure 4.15a, 4.17a, 4.19a, 4.21a, and 4.23a for the five controllers, respectively. These results indicate that trajectory termination is primarily caused by violation of the roll angle φ constraint. This behavior is observed for all controllers except MRAC, which achieves the highest success rate. However, this improvement comes at the expense of tracking performance, as MRAC also yields the highest RMSE across both the pitch and yaw channels, as reported in Table 4.4–4.6. In contrast, the \mathcal{L}_1 CMRAC controllers share the lowest RMSE, indicating improved acceleration tracking performance, albeit with a marginally lower success rate.

Path 3

The performance metrics for Path 3 broadly follow the trends observed for Path 2. In particular, the conservative \mathcal{L}_1 CMRAC outperforms both the nominal and aggressive \mathcal{L}_1 CMRAC configurations, as well as the LQ controller, in terms of success rate. A notable deviation, however, is the performance of the MRAC, which exhibits the poorest results across all metrics in both the pitch and yaw channels. As shown in Table 5.1, the success rate is reduced by more than 30% compared to the best-performing controller, namely the conservative \mathcal{L}_1 CMRAC.

Inspection of the MC traces for MRAC reveals that the primary cause of unsuccessful simulations is violation of the lateral acceleration constraint a_y , as illustrated in Figure 4.26. This indicates a limited ability of the controller to regulate the yaw channel under the increased excitation present in Path 3.

Furthermore, a comparison of the RMSE across all paths in the pitch channel (Tables 4.2, 4.5, and 4.8) shows that MRAC consistently exhibits inferior tracking performance in vertical acceleration. This behavior is particularly evident in Path 2, where Figure 4.16b shows persistent overshoot of the soft constraint limit at approximately 200 m/s^2 , a trend also observed in the nominal case.

These results suggest that the MRAC is not optimally tuned for the considered operating conditions. In combination with the increased cross-coupling arising from simultaneous excitation of both pitch and yaw channels, this leads to degraded performance and a significantly reduced success rate in Path 3 when large parameter uncertainties are present.

5.4 External Disturbance Attenuation

When the vehicle is subjected to an external disturbance, both the LQ controller and the conservative \mathcal{L}_1 CMRAC demonstrate comparable disturbance attenuation, whereas the MRAC fails to maintain stability. The results are illustrated in Figure 4.35, where it can be observed that the lateral acceleration for the MRAC exceeds the imposed constraint when the disturbance is removed at $t = 25 \text{ s}$. This behavior is consistent with the results presented in Section 5.3.1, where the MRAC exhibited the poorest tracking performance among the evaluated controllers, particularly in the pitch and yaw channels. The reduced tracking accuracy limits its ability to handle large disturbances, ultimately leading to constraint violations.

Comparing the two successful controllers, the LQ and the conservative \mathcal{L}_1 CMRAC exhibit similar overall performance. However, the \mathcal{L}_1 CMRAC provides superior performance in the roll and yaw channels, where it achieves both reduced overshoot and faster settling compared to the LQ controller. In contrast, in the pitch channel, the \mathcal{L}_1 CMRAC exhibits slightly increased overshoot while maintaining a settling time comparable to that of the LQ controller, particularly around $t = 22$ s when the disturbance is introduced.

The adaptive weights, shown in Figure 4.34, further highlight the differences between the adaptive controllers. The MRAC responds to the disturbance with significant activity across multiple adaptive weights, indicating a more aggressive and distributed adaptation. In contrast, the \mathcal{L}_1 CMRAC exhibits a more structured response involving fewer active weights, reflecting the effect of the filtered adaptation. Notably, the magnitude of the active weights is larger, which suggests that the adaptation is concentrated in a smaller subset of parameters rather than distributed across many, as observed in the MRAC.

It is also noted that the variation in the \mathcal{L}_1 CMRAC weights begins prior to the disturbance onset. As seen in Figure 4.34b, this behavior is likely associated with the transient in vertical acceleration observed before $t = 22$ s in Figure 4.35c, rather than the external disturbance itself.

5.5 Accuracy–Robustness Trade-off

The results reveal a fundamental tension between tracking performance and robustness that is inherent to the uncertainty bound and predictor design within the \mathcal{L}_1 CMRAC framework. This trade-off, whereby improved tracking accuracy comes at the cost of reduced robustness margins, is a well-established principle in control theory and manifests clearly across the evaluated trajectories. Tightening the uncertainty bound to enforce stricter satisfaction of the \mathcal{L}_1 small-gain condition in (3.27) necessitates a more aggressive predictor, which introduces competing effects whose net impact depends on the local characteristics of the flight envelope traversed.

As illustrated in Figure 4.2, the tightly tuned uncertainty bound drives a high-bandwidth predictor that satisfies the \mathcal{L}_1 small-gain condition at the designated operating points in Figure 3.3 but produces a locally elevated and non-uniform \mathcal{L}_1 norm between the operating points, due to amplification of model mismatch introduced by linearization and gain-scheduling interpolation. As characterized in Section 5.2, this elevation directly determines the robustness margins of the controller, with a larger \mathcal{L}_1 norm corresponding to reduced margins and increased risk of instability. The impact of this elevation on closed-loop performance is observed to depend strongly on which region of the flight envelope the trajectory traverses, as determined by the Mach profiles shown in Figures 3.7c, 3.8c, and 3.9c.

For Path 1, the Mach profile in Figure 3.7c indicates that the trajectory operates in a region of the flight envelope where the aggressive controller's \mathcal{L}_1 norm remains below the small-gain bound, as shown in Figure 4.2. In this region, the less aggressive predictor of the conservative controller produces a norm that exceeds the small-gain bound, prioritizing tracking performance at the cost of robustness margins. Consistent with this interpretation, the conservative controller achieves lower

RMSE but a reduced success rate compared to the aggressive variant (see Tables 4.1 and 4.3). The aggressive controller, operating below the small-gain bound, exhibits increased robustness margins at the cost of tracking precision, yielding higher RMSE but an improved success rate. This behavior is consistent with \mathcal{L}_1 theory, which predicts that a more restrictive filter on the adaptive control signal reduces tracking aggressiveness while increasing stability margins.

For Path 2, the Mach profile in Figure 3.8c indicates that the trajectory traverses a transitional region of the flight envelope. Accordingly, the performance difference between the two controllers is negligible across the evaluated metrics (see Tables 4.4–4.6), and no conclusive assessment of the accuracy–robustness trade-off can be drawn from this path alone.

The Mach profile for Path 3, illustrated in Figure 3.9c, implies that the trajectory operates in a region where the aggressive controller’s \mathcal{L}_1 norm is most elevated, as shown in Figure 4.2. From \mathcal{L}_1 theory, a larger norm corresponds to a less restrictive filter on the adaptive control signal, permitting more aggressive tracking action [2]. Consistent with this, the aggressive controller achieves lower RMSE than the conservative variant (see Tables 4.8 and 4.9). However, operating closer to the stability margin associated with the elevated norm increases the risk of failure, resulting in a reduced success rate. Conversely, the conservative controller, whose lower-bandwidth predictor does not amplify model mismatch to the same degree, operates with a norm significantly closer to the bound than the aggressive variant, as shown in Figure 4.3. This yields higher RMSE but stronger robustness margins, resulting in a higher success rate. The consistency of this behavior with the accuracy–robustness trade-off across both controllers and both metrics strengthens the interpretation that the \mathcal{L}_1 norm elevation is the primary determinant of the balance between tracking precision and robustness in this application.

These observations collectively highlight that violation of the small-gain condition does not preclude stable and competitive closed-loop performance, as evidenced by Figure 4.3 and Table 5.1. This is consistent with \mathcal{L}_1 theory, which establishes the small-gain condition as a sufficient condition for stability and bounded performance, not a necessary one [2]. Despite this, the controller maintains stable closed-loop behavior throughout, indicating that violation implies reduced robustness margins rather than instability, as characterized in Section 5.2. This is corroborated by Table 5.1, which shows that the conservative \mathcal{L}_1 CMRAC consistently exceeds the success rate of the LQ controller across all evaluated paths. Since the LQ controller operates without any adaptive robustness mechanism, this comparison provides empirical evidence that the adaptive architecture retains meaningful robustness benefits even when the small-gain condition is not strictly satisfied. Furthermore, the aggressive controller, which exhibits a significantly more elevated \mathcal{L}_1 norm in the regions traversed by Path 3 as shown in Figure 4.2, achieves a success rate within 1% of the conservative variant for that path, as reported in Table 5.1. This further supports the interpretation that the small-gain condition is conservative in practice, and that moderate violation does not deterministically lead to instability or significant performance degradation.

Furthermore, crashes are attributed to constraint violation driven by residual model mismatch and parameter uncertainty that the adaptive law fails to fully compen-

sate, rather than to instability induced by small-gain condition violation. The \mathcal{L}_1 small-gain condition in (3.27) ensures only that the filtered adaptive control signal does not destabilize the prediction error dynamics, and provides no guarantee that residual unestimated uncertainty remains sufficiently small to prevent constraint violation through degraded tracking. Consequently, even the conservative controller experiences some degree of performance degradation and occasional constraint violation, though to a lesser extent than the aggressive variant in the elevated norm regions.

As demonstrated in Section 5.2, strict satisfaction of the \mathcal{L}_1 small-gain condition at discrete operating points does not guarantee its satisfaction across the flight envelope when gain scheduling is employed, and the resulting locally elevated and non-uniform \mathcal{L}_1 norm directly governs the robustness margins available to the controller along each trajectory. A moderate relaxation of the small-gain condition, achieved by loosening the uncertainty bound and reducing predictor bandwidth, yields a flatter norm profile and more consistent performance across the flight envelope. As the primary mechanism through which the uncertainty bound influences closed-loop behavior, the predictor design constitutes the most critical tuning choice in the \mathcal{L}_1 CMRAC framework.

Beyond the trade-off between the two \mathcal{L}_1 CMRAC variants, the results further demonstrate that a purely performance-oriented adaptive controller does not necessarily yield improved performance. The MRAC controller, which is tuned primarily for tracking performance, exhibits the highest RMSE in all scenarios (see Tables 4.1–4.3, 4.4–4.6, and 4.7–4.9) and performs particularly poorly on Path 3, where it yields a significantly reduced success rate of 60% whereas the \mathcal{L}_1 CMRAC tuned for a small uncertainty bound gave above 92%. These results indicate that prioritizing performance alone, without accounting for robustness, can degrade overall controller effectiveness.

This ability of the \mathcal{L}_1 CMRAC to operate in a regime where robustness and performance can be jointly tuned is fundamentally enabled by the composite structure of the \mathcal{L}_1 CMRAC. In particular, the outer-loop is responsible for enforcing tracking performance and penalize deviations from the nominal reference model, while the inner-loop ensures robustness through rapid estimation and filtering of destabilizing uncertainties. This structural separation yields a control architecture that allows a partial decoupling of performance and robustness objectives.

Taken together, the results suggest that the \mathcal{L}_1 CMRAC should be tuned to achieve sufficiently fast predictor dynamics, while avoiding overly aggressive designs that amplify model mismatch. Although this may lead to slight violations of the \mathcal{L}_1 small-gain condition, maintaining the predictor dynamics within a reasonable range results in improved practical performance. Thus, the \mathcal{L}_1 CMRAC enables a design that balances robustness and tracking performance through its composite structure, rather than strictly prioritizing either extreme.

5.6 Practical Comparison

While the \mathcal{L}_1 CMRAC achieves a slightly higher success rate than the LQ controller across the three paths, the difference in overall performance is relatively small con-

sidering the increased complexity of the adaptive architecture. This indicates that the simpler LQ controller remains competitive in terms of practical performance, despite not incorporating adaptation.

In contrast, the \mathcal{L}_1 CMRAC requires a significantly larger number of design choices, including multiple tuning parameters associated with the predictor dynamics and adaptive law. This increases the complexity of the controller design process compared to the LQ approach.

Furthermore, the \mathcal{L}_1 CMRAC introduces increased computational complexity due to the evaluation of terms such as $\hat{\Omega}^\top \Phi(x)$ in the adaptive update law, resulting in a higher online computational burden.

Consequently, a trade-off arises between the modest performance improvement achieved by the \mathcal{L}_1 CMRAC and the increased design and computational complexity. Importantly, the results also highlight that the LQ controller, despite its simplicity and lack of adaptation, remains close in performance to the adaptive solution in all evaluated scenarios.

6

Conclusion

To conclude this work, the research questions provided in Section 1.4 are revisited. After the questions are answered, future work is presented.

6.1 Question 1: Robustness Guarantees

The first research question asks: What robustness guarantees can be provided for the designed adaptive controller? As presented in Section 4 and discussed in Section 5, the \mathcal{L}_1 CMRAC architecture guarantees boundedness of both the prediction error and the adaptive parameter estimates via the \mathcal{L}_1 small-gain condition in 3.27, and the projection operator described in Section 3.2.3. These properties ensure that the adaptive updates remain well-behaved and that the predictor dynamics do not lead to instability, even under fast adaptation.

The robustness of the \mathcal{L}_1 CMRAC is formally characterized through the \mathcal{L}_1 small-gain condition in (3.27). However, the results demonstrate that this condition is conservative. In particular, violation of the small-gain condition does not directly imply instability of the closed-loop system. Instead, it reflects a loss of theoretical guarantees rather than a loss of practical stability. This is supported by the observed performance, where the conservative \mathcal{L}_1 CMRAC, despite violating (3.27) across the operating points in Figure 3.3, achieves the highest overall success rate among all evaluated controllers.

Conversely, the aggressive \mathcal{L}_1 CMRAC, which satisfies the small-gain condition through faster predictor dynamics, exhibits reduced tracking performance and the lowest success rate among the \mathcal{L}_1 -based controllers. This highlights an inherent trade-off: enforcing the small-gain condition requires fast predictor dynamics, which amplifies model mismatch and degrades performance.

In comparison to the LQ controller, the \mathcal{L}_1 CMRAC demonstrates improved robustness, as reflected by its consistently higher success rates under uncertainty. However, this improvement comes at the cost of increased design complexity and computational burden, as discussed in Section 5.6.

In summary, the \mathcal{L}_1 CMRAC provides stronger robustness guarantees in terms of bounded adaptive behavior and improved empirical performance under uncertainty. Nevertheless, these advantages must be weighed against the increased implementation complexity, and the conservative nature of the small-gain condition should be carefully considered in practical tuning.

6.2 Question 2: Balance Between Robustness and Performance

The second research question asks: Is it necessary, even with adaptive control, to balance robustness against performance? The results demonstrate that such a trade-off is indeed fundamental, even within adaptive control frameworks.

A purely performance-oriented adaptive controller, such as the MRAC, is not able to consistently attenuate uncertainties, which leads to degraded performance and reduced success rates, specifically in Path 3. In contrast, overly robust \mathcal{L}_1 CMRAC configurations, achieved through aggressive predictor dynamics to strictly satisfy the \mathcal{L}_1 small-gain condition, result in reduced tracking performance and an increased likelihood of constraint violations, specifically due to model mismatch introduced by gain-scheduling.

The results show that neither extreme is desirable. Instead, the best overall performance is obtained with a moderately tuned \mathcal{L}_1 CMRAC, where robustness is enforced without excessively increasing the closed-loop predictor bandwidth. In particular, the conservative \mathcal{L}_1 CMRAC achieves the highest overall success rate, demonstrating that a balanced design provides the most favorable trade-off between robustness and tracking performance.

An additional key observation is that the cascade structure of the \mathcal{L}_1 CMRAC plays an important role in enabling this trade-off. The architecture uses the tracking error in the outer-loop adaptive law to drive performance, while the inner-loop predictor and low-pass filter are driven by the prediction error to ensure bounded parameter updates and filtered uncertainty compensation. This structured use of different error signals introduces a functional decoupling between performance shaping and robustness enforcement. As a result, the controller allows a practical balance between tracking performance and robustness, rather than enforcing one at the expense of the other.

Therefore, it is concluded that adaptive control does not eliminate the classical trade-off between robustness and performance. Rather, it shifts the design problem toward finding an appropriate balance between these objectives, where both insufficient robustness and excessive conservatism in the \mathcal{L}_1 small-gain condition can degrade overall system performance.

6.3 Future Work

The following directions are proposed for further evaluation and potential improvement of the \mathcal{L}_1 CMRAC framework with respect to robustness and performance.

- Increase the number of MC simulations to obtain higher statistical confidence in the reported success rates and performance metrics.
- Evaluate the controllers on a wider set of reference trajectories in order to assess performance across a broader range of operating conditions and maneuver types.
- Perform a more comprehensive sensitivity analysis in the frequency domain to better characterize the robustness margins and to identify frequency ranges

where model mismatch and interpolation effects dominate system behavior.

- Replace the outer-loop controller of the \mathcal{L}_1 CMRAC with a model predictive controller (MPC) [22] to explicitly enforce state and input constraints and reduce the risk of constraint violations.
- Implement a scheduling-law which adjusts the aggressiveness of the closed-loop predictor dynamics depending on the magnitude of the interpolation error. Thus, the elevated \mathcal{L}_1 norms in shown in Figure 4.2 might be mitigated.
- Remove decoupling of the control channels and apply the \mathcal{L}_1 CMRAC on the multiple input multiple output (MIMO) system. Thus, the cross-couplings are explicitly described in the model to improve uncertainty attenuation.

Bibliography

- [1] M. Lourié and G. Carefall, “Adaptive control of high-speed aerial vehicles using mrac methods,” Master’s thesis, Linköping University, 2025. [Online]. Available: <https://www.diva-portal.org/smash/get/diva2:1973027/FULLTEXT01.pdf>
- [2] N. Hovakimyan, C. Cao, E. Kharisov, E. Xargay, and I. M. Gregory, “L1 adaptive control for safety-critical systems,” Purdue University, Tech. Rep., 2008. [Online]. Available: https://engineering.purdue.edu/~zak/ECE675__2022/L1_adaptive_control.pdf
- [3] SAAB AB, “Code of conduct,” 2023. [Online]. Available: https://www.saab.com/globalassets/corporate/sustainability/coc/saab_coc_en_230331b.pdf
- [4] B. Kulcsar, “Robust and nonlinear control,” 2024, lecture notes for EEN050 Robust and Nonlinear Control, Systems and Control Department, Chalmers University of Technology.
- [5] PlanetMath.org, “Hurwitz matrix,” 2013, accessed: 2026-02-25. [Online]. Available: <https://planetmath.org/hurwitzmatrix>
- [6] T. Glad and L. Ljung, *Control Theory: Multivariable and Nonlinear Methods*. London: Taylor & Francis, 2000.
- [7] M. Sato and C. Edwards, “Robustness against matched uncertainty by linear state-feedback controller combined with internal saturation,” *IFAC-PapersOnLine*, vol. 59, no. 16, pp. 349–354, 2025.
- [8] E. Lavretsky and K. A. Wise, *Robust and Adaptive Control*, 2nd ed. Cham, Switzerland: Springer, 2024.
- [9] J. C. A. Barata and M. S. Hussein, “The moore–penrose pseudoinverse: A tutorial review of the theory,” *arXiv preprint arXiv:1110.6882*, 2011. [Online]. Available: <https://arxiv.org/abs/1110.6882>
- [10] G. Chowdhary and E. N. Johnson, “Concurrent learning for convergence in adaptive control without persistency of excitation,” in *Proceedings of the 49th IEEE Conference on Decision and Control (CDC)*, Atlanta, GA, USA, 2010, pp. 3674–3679.
- [11] A. Megretski, “Small gain theorem lecture notes,” 2006, 6.241 Dynamic Systems and Control, Massachusetts Institute of Technology, Fall 2006. [Online]. Available: https://web.mit.edu/6.241/ameg_www_fall2006/www/images/L06smallgain.pdf
- [12] MathWorks. (2023) Low-pass filter. Accessed: 2026-03-02. [Online]. Available: <https://se.mathworks.com/discovery/low-pass-filter.html>
- [13] ScienceDirect. (2023) Cutoff frequency. Accessed: 2026-03-02. [Online]. Available: <https://www.sciencedirect.com/topics/computer-science/cutoff-frequency>

- [14] MathWorks, “Gain scheduled control systems,” MATLAB & Simulink Control Design, 2026, online; accessed 2026-02-27. [Online]. Available: <https://se.mathworks.com/help/control/ug/gain-scheduled-control-systems.html>
- [15] NASA Glenn Research Center, “Speed of sound,” n.d., accessed: 2026-04-23. [Online]. Available: <https://www.grc.nasa.gov/www/k-12/VirtualAero/BottleRocket/airplane/sound.html>
- [16] “Cascade control,” ScienceDirect Topics, 2023, accessed: 2026-03-03. [Online]. Available: <https://www.sciencedirect.com/topics/engineering/cascade-control>
- [17] X. Lu, Q. Geng, and R. Mo, “New composite mrac with modification for fixed-wing uav,” in *Proceedings of the Chinese Control and Decision Conference (CCDC)*. IEEE, 2016, pp. 479–484, paper No. 7531032. [Online]. Available: <https://ieeexplore.ieee.org/document/7531032>
- [18] H. Saoud, “Composite lyapunov criteria for stability and convergence with applications to optimization dynamics,” arXiv preprint, 2025, arXiv:2510.08259. [Online]. Available: <https://arxiv.org/abs/2510.08259>
- [19] ScienceDirect, “Monte carlo simulation,” <https://www.sciencedirect.com/topics/physics-and-astronomy/monte-carlo-simulation>, n.d., accessed: 2026-03-26.
- [20] L. Wasserman, *All of Statistics: A Concise Course in Statistical Inference*, ser. Springer Texts in Statistics. Springer, 2004.
- [21] Swedish Meteorological and Hydrological Institute (SMHI), “Wind speed scales,” n.d., accessed: 2026-04-22. [Online]. Available: <https://www.smhi.se/kunskapsbanken/meteorologi/vind/skalar-for-vindhastighet>
- [22] N. Murgovski, “Model predictive control,” Lecture notes for SSY281 Model Predictive Control, Systems and Control Department, 2024.

A

Appendix 1

A.1 Evaluation of the \mathcal{L}_1 CMRAC Architecture Using a Constant-Velocity Model

To provide an initial validation of the proposed \mathcal{L}_1 CMRAC architecture, a simplified one-dimensional constant-velocity model is considered. The discrete-time system is given by

$$\begin{aligned}x_{k+1} &= \begin{bmatrix} 1 & h \\ 0 & 1 \end{bmatrix} x_k + \begin{bmatrix} 0 \\ h \end{bmatrix} u_k, \\ y_k &= \begin{bmatrix} 1 & 0 \end{bmatrix} x_k,\end{aligned}\tag{A.1}$$

where $h = 0.1$ denotes the sampling time. The state vector is defined as

$$x_k = \begin{bmatrix} p_k & v_k \end{bmatrix}^\top,\tag{A.2}$$

where p represents the position and v the velocity. The control input u corresponds to the applied acceleration.

The controllers are evaluated on the reference trajectory illustrated in Figure A.1a, under bounded parametric uncertainties characterized by

$$\Omega_{\max} = 3.431, \quad \Lambda_{\max} = 1.334.\tag{A.3}$$

Both an LQ controller and the proposed \mathcal{L}_1 CMRAC controller are evaluated over 1000 Monte Carlo simulations. The resulting trajectories are shown in Figure A.1b–A.1c, providing a qualitative comparison of tracking performance and robustness under uncertainty. Success is based on the ability to stay within a predetermined area, more specifically

$$\begin{aligned}-10 &\leq p_k \leq 10 \\ -6 &\leq v_k \leq 6\end{aligned}\tag{A.4}$$

As Figure A.1 shows, the \mathcal{L}_1 CMRAC outperforms the LQ by close to 50% and a reduced RMSE by approximately factor $10^{-2} \times$. These results show the potential of the \mathcal{L}_1 CMRAC, where due to the cascade structure, the controller is able to maintain stability under large parameter uncertainties, without sacrificing performance accuracy.

Table A.1: Performance results of each controller for position. RMSE and MAE are averaged over all MC simulations.

Controller	MAE	RMSE	Success Rate
LQR	2.656	1.815	51.3%
\mathcal{L}_1 CMRAC	0.133	0.080	100%

Table A.2: Performance results of each controller for velocity. RMSE and MAE are averaged over all MC simulations.

Controller	MAE	RMSE	Success Rate
LQR	2.336	1.573	51.3%
\mathcal{L}_1 CMRAC	0.280	0.012	100%

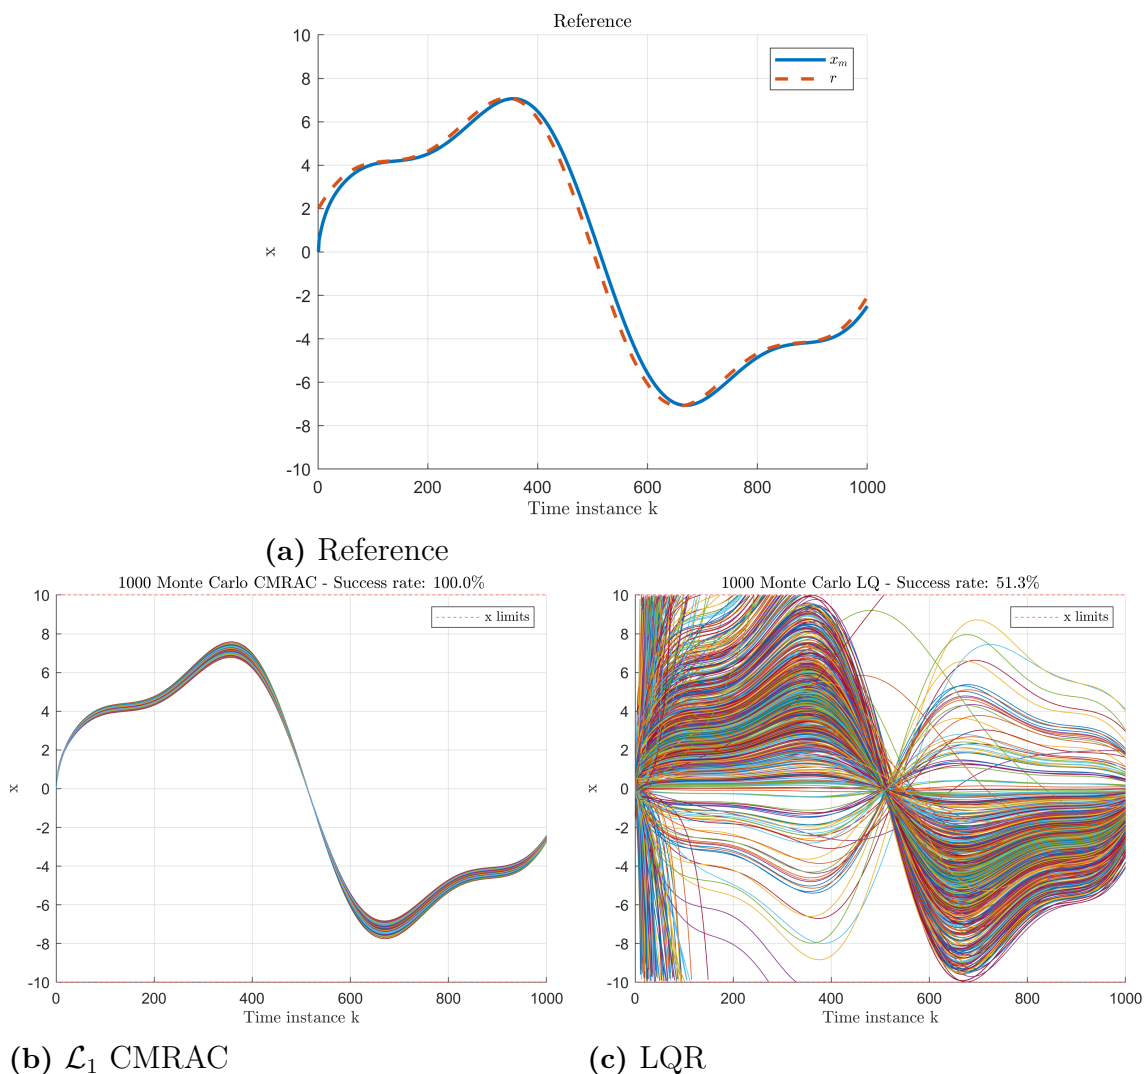


Figure A.1: Monte Carlo trajectory realizations for the \mathcal{L}_1 CMRAC and LQ controllers applied to the 1D constant-velocity model.

DEPARTMENT OF ELECTRICAL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden
www.chalmers.se



CHALMERS
UNIVERSITY OF TECHNOLOGY