



FEDERAL UNIVERSITY OF PARÁ
INSTITUTE OF TECHNOLOGY
POSTGRADUATE PROGRAM IN ELECTRICAL ENGINEERING

Robust Control for Aerodynamic Steering of a Reusable Launch Vehicle: An LQG/LTR Approach with Wind Disturbance Rejection

Rufin Azonsivo

DM 32/2025

UFPA / ITEC / PPGEE
Guamá University Campus
Belém-Pará-Brasil

2025

FEDERAL UNIVERSITY OF PARÁ
INSTITUTE OF TECHNOLOGY
POSTGRADUATE PROGRAM IN ELECTRICAL ENGINEERING

Rufin Azonsivo

**Robust Control for Aerodynamic Steering of a Reusable Launch
Vehicle: An LQG/LTR Approach with Wind Disturbance Rejection**

Master's dissertation submitted to the
Examination Committee of the Graduate
Program in Electrical Engineering at UFPA
for the fulfillment of the Master's Degree in
Electrical Engineering in the field of Control
and Automation.

Advisor: Prof. Dr. Antonio da Silva Silveira

Co-advisor: Prof. Dr. Marcos Cesar da Rocha Seruffo

UFPA / ITEC / PPGEE
Guamá University Campus
Belém-Pará-Brasil

2025

**Dados Internacionais de Catalogação na Publicação (CIP) de acordo com ISBD
Biblioteca do Instituto de Tecnologia / UFPA – Belém-PA**

A996r Azonsivo, Rufin

Robust control for aerodynamic steering of a reusable launch vehicle: an LQG/LTR approach with Wind Disturbance Rejection / Rufin Azonsivo . — 2025.

Orientador: Antonio da Silva Silveira

Coorientador: Marcos Cesar da Rocha Seruffo

Dissertação (Mestrado) – Universidade Federal do Pará, Instituto de Tecnologia, Programa de Pós-Graduação em Engenharia em Engenharia Elétrica, Belém, 2025.

1. Controle robusto 2. Controle automático 3. Aerodinâmica I. Título.

CDD 23. ed. – 629.8


Elaborado por Maria Ivone Maia da Costa – CRB-2/768




UNIVERSIDADE FEDERAL DO PARÁ
INSTITUTO DE TECNOLOGIA
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

ATA DA APRESENTAÇÃO E DEFESA DE DISSERTAÇÃO DE MESTRADO

ATA DA 832ª SESSÃO DE APRESENTAÇÃO DE DEFESA DE DISSERTAÇÃO DE MESTRADO PARA CONCESSÃO DE GRAU DE MESTRE EM ENGENHARIA ELÉTRICA NA ÁREA DE **SISTEMAS DE ENERGIA ELÉTRICA** REALIZADA ÀS QUATORZE HORAS DO DIA DEZ DE OUTUBRO DE DOIS MIL E VINTE E CINCO, VIA PLATAFORMA GOOGLE MEET, INTITULADA: **ROBUST CONTROL FOR AERODYNAMIC STEERING OF A REUSABLE LAUNCH VEHICLE: AN LQG/LTR APPROACH WITH WIND DISTURBANCE REJECTION** APRESENTADA DURANTE 70 MINUTOS PELO CANDIDATO **RUFIN AZONSIVO** DIANTE DA BANCA EXAMINADORA APROVADA PELO COLEGIADO DO PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA DO INSTITUTO DE TECNOLOGIA DA UNIVERSIDADE FEDERAL DO PARÁ, ASSIM CONSTITUÍDA: **PROF. DR. ANTONIO DA SILVA SILVEIRA** (ORIENTADOR - PPGE/ITEC/UFPA), **PROF. DR. MARCOS CÉSAR DA ROCHA SERUFFO** (COORIENTADOR - PPGE/ITEC/UFPA), **PROF. DR. WALTER BARRA JÚNIOR** (AVALIADOR INTERNO - PPGE/ITEC/UFPA), **PROF. DR. FRANCISCO DAS CHAGAS DE SOUZA** (AVALIADOR EXTERNO - UFMA), **PROF. DR. JOÃO VIANA DA FONSECA NETO** (AVALIADOR EXTERNO - UFMA) E **PROF. DR. JOSÉ ALANO PERES DE ABREU** (AVALIADOR EXTERNO - CLA). CONCLUÍDOS OS TRABALHOS DE APRESENTAÇÃO E ARGUIÇÃO, A BANCA EXAMINADORA DECIDIU PELA **APROVAÇÃO** DO CANDIDATO. FOI CONCEDIDO UM PRAZO DE TRINTA DIAS, PARA O CANDIDATO EFETUAR AS CORREÇÕES SUGERIDAS PELA COMISSÃO EXAMINADORA E APRESENTAR O TRABALHO EM SUA REDAÇÃO DEFINITIVA. ESTA ATA NÃO VALE COMO OUTORGA DE GRAU DE MESTRADO, DE ACORDO COM O DEFINIDO NA RESOLUÇÃO 072/2004-CONSEPE. E, PARA CONSTAR, FOI LAVRADA A PRESENTE ATA, QUE VAI ASSINADA PELOS MEMBROS DA COMISSÃO E PELO CANDIDATO.

Documento assinado digitalmente
 **ANTONIO DA SILVA SILVEIRA**
Data: 14/10/2025 15:56:56-0300
Verifique em <https://validar.iti.gov.br>

Prof. Dr. Antonio da Silva Silveira
(Orientador – PPGE/ITEC/UFPA)

Documento assinado digitalmente
 **MARCOS CESAR DA ROCHA SERUFFO**
Data: 14/10/2025 18:41:23-0300
Verifique em <https://validar.iti.gov.br>

Prof. Dr. Marcos César da Rocha Seruffo
(Coorientador - PPGE/ITEC/UFPA)

() **VIDE VERSO:** Em caso de alteração do título pela Comissão Examinadora.



UNIVERSIDADE FEDERAL DO PARÁ
INSTITUTO DE TECNOLOGIA
PROGRAMA DE PÓS-GRADUAÇÃO EM ENGENHARIA ELÉTRICA

Documento assinado digitalmente



WALTER BARRA JÚNIOR
Data: 15/10/2025 16:05:53-0300
Verifique em <https://validar.iti.gov.br>

Prof. Dr. Walter Barra Júnior
(Avaliador Interno – PPGEE/ITEC/UFPA)

Documento assinado digitalmente



FRANCISCO DAS CHAGAS DE SOUZA
Data: 14/10/2025 16:54:54-0300
Verifique em <https://validar.iti.gov.br>

Prof. Dr. Francisco das Chagas de Souza
(Avaliador Externo – UFMA)

Documento assinado digitalmente



JOAO VIANA DA FONSECA NETO
Data: 14/10/2025 19:41:06-0300
Verifique em <https://validar.iti.gov.br>

Prof. Dr. João Viana da Fonseca Neto
(Avaliador Externo – UFMA)

Documento assinado digitalmente



JOSE ALANO PERES DE ABREU
Data: 14/10/2025 21:26:10-0300
Verifique em <https://validar.iti.gov.br>

Prof. Dr. José Alano Peres de Abreu
(Avaliador Externo – CLA)

Documento assinado digitalmente



RUFIN AZONSIVO
Data: 14/10/2025 17:07:30-0300
Verifique em <https://validar.iti.gov.br>

CANDIDATO: _____

Por sugestão da Comissão Examinadora, o novo título passa a ser:

Documento assinado digitalmente



ANTONIO DA SILVA SILVEIRA
Data: 14/10/2025 15:58:19-0300
Verifique em <https://validar.iti.gov.br>

Prof. Dr. Antonio da Silva Silveira
(Orientador – PPGEE/ITEC/UFPA)

() **VIDE VERSO:** Em caso de alteração do título pela Comissão Examinadora.

I dedicate this work to the cherished memory of my beloved father, Gaston, and my dearly departed brothers, Omer and Charly. Though absent from this world, their presence remains eternally engraved in my heart.

Acknowledgements

First and foremost, I thank God for the strength, health, courage, and determination granted to me, which were essential for the completion of this work. Without Him, none of this would have been possible.

I am deeply grateful to my mother Josephine, my uncle Lazare, and to my siblings, Saturnin, Clémentine, Tanguy, Rufine and Chardace, for their unwavering support and encouragement throughout this journey. I also extend my heartfelt thanks to my cherished children, Fifamè Arielle and Fênou Arthur, for always brightening my days.

I owe immense gratitude to my advisor, Professor Dr. Antonio da Silva Silveira, for his invaluable teachings, exceptional guidance, friendship, patience, and belief in my potential during the development of this work.

I would like to thank my co-advisor, Professor Dr. Marcos Cesar da Rocha Seruffo, for his encouragement invaluable helping during my master's studies.

My sincere thanks go to all my friends in the Postgraduate Program in Electrical Engineering at the Institute of Technology (ITEC), especially Victor, Matheus, Elder, Lucas, Daniel, and Marcelo for the insightful discussions, support, and motivation during the years of this work. I also thank all the friends I made at UFPA during my master's coursework, Moïses, Expédit, Geraldine, Camille, Alphonse, Marius, Borgia for their various contributions, comforts and encouragement.

I am grateful to the Federal University of Pará (UFPA), particularly the Institute of Technology through the Postgraduate Program in Electrical Engineering (PPGEE) for their institutional support in the development of this research.

I acknowledge the Coordination for the Improvement of Higher Education Personnel (CAPES) for their financial support.

To all those who, in any way, contributed to the realization of this work thank you.

“Between stimulus and response, there is a space. In that space is our power to choose our response. In our response lies our growth and our freedom.”
(Viktor Frankl, Man’s Search for Meaning)

Abstract

Facing rising costs in space exploration, the aerospace industry has embraced reusability as a transformative solution. A paradigm shift pioneered by SpaceX in 2015 and now adopted globally. This innovation introduces multidisciplinary complexities, particularly in controlling reusable boosters and first stages during atmospheric re-entry and descent. Precise attitude control and aerodynamic manoeuvring are critical to guide these components along predetermined trajectories while ensuring stable deceleration. By leveraging advanced control systems and aerodynamic design, modern launch vehicles achieve recovery and reuse, dramatically reducing mission costs. This dissertation focuses on attitude dynamics and explores a Linear Quadratic Gaussian (LQG) solution for the controlled aerodynamic descent of a reusable launcher using the Loop Transfer Recovery (LTR) method. This technique is based on the fact that the Linear Quadratic Regulator (LQR) using state-variable feedback has specific guaranteed robustness properties. The LQG, like the association of LQR and the Kalman Filter (KF), is tested on a decoupled mathematical model of a launcher vehicle based on the prerequisites of the reusable rocket demonstrator CALLISTO, Cooperative Action Leading to Launcher Innovation in Stage Toss-back Operations, and the existing literature in the field. Using LQG/LTR as an optimal and stochastic control association to propose a robust control of position by direct action on the whole attitude of the vehicle body described by Euler angles, this work aims to investigate the stability and robustness performances of the latter for the control design of the aerodynamic descent phase of a reusable rocket. Therefore, the control structure designed on different cases and the possibility of in-flight events of the rocket in its controlled aerodynamic descent are discussed. Also, different correction strategies are used to refine the initially set performance objectives, like the predictive and weighted guidance trajectory.

Key-words: Reusable launch vehicle, position control, rocket dynamics, aircraft control, LQG/LTR

Resumo

Diante dos crescentes custos na exploração espacial, a indústria aeroespacial adotou a reutilização como solução transformadora. Uma mudança de paradigma iniciada pela SpaceX em 2015 e agora difundida globalmente. Essa inovação introduz complexidades multidisciplinares, especialmente no controle de boosters e primeiros estágios reutilizáveis durante a reentrada atmosférica e descida. Controle preciso de atitude e manobras aerodinâmicas são críticos para guiar esses componentes ao longo de trajetórias predeterminadas, garantindo desaceleração estável. Ao empregar sistemas de controle avançados e projetos aerodinâmicos, os veículos lançadores modernos alcançam recuperação e reuso, reduzindo drasticamente os custos das missões. Esta dissertação foca na dinâmica de atitude e explora uma solução Linear Quadratic Gaussian (LQG) para a descida aerodinâmica controlada de um veículo lançador reutilizável, utilizando o método Loop Transfer Recovery (LTR). Essa técnica baseia-se no fato de que o regulador linear quadrático (LQR) com realimentação de estados possui propriedades de robustez garantidas. O LQG, como a associação entre o LQR e o filtro de Kalman (KF), é testado em um modelo matemático desacoplado de um veículo lançador, baseado nos requisitos do demonstrador reutilizável CALLISTO (Cooperative Action Leading to Launcher Innovation in Stage Toss-back Operations) e na literatura existente. Ao usar LQG/LTR como uma associação de controle ótimo e estocástico, este trabalho propõe um controle robusto de posição por ação direta na atitude global do veículo (descrita por ângulos de Euler), investigando sua estabilidade e robustez para o projeto de controle da fase de descida aerodinâmica de um foguete reutilizável. A estrutura de controle é projetada para diferentes cenários, incluindo eventos em voo durante a descida controlada. Além disso, estratégias de correção, como trajetórias preditivas e ponderadas, são aplicadas para refinar os objetivos de desempenho iniciais.

Palavras-chave: Veículo lançador reutilizável, controle de posição, dinâmica de foguetes, controle aeronáutico, LQG/LTR

List of figures

Figure 2.1 – Mission illustration in left and 3D plotting of a suggested trajectory for back to launch site scenario by right.	11
Figure 2.2 – Inertial UEN and Downrange–Crossrange–Altitude (DCA) references frame emphasizing \mathcal{X}_T rotation to decoupling.	12
Figure 2.3 – Body references frame.	12
Figure 2.4 – 3-2-1 Rotation sequence to move from Downrange–Crossrange–Altitude to body axes.	13
Figure 2.5 – Representation of body reference frame, air-path reference frame, angle of attack, and sideslip angle.	13
Figure 2.6 – Illustration of scheduled points and two interpolation laws for a scalar gain $K(\rho)$: piecewise linear (dashed) and cubic spline (solid).	24
Figure 3.1 – Proposed operation points around a suggested trajectory for back to lunch site scenario by right.	28
Figure 3.2 – CALLISTO-based RLV control loop for the aerodynamic descent subject to wind disturbance.	32
Figure 3.3 – Block diagram of the LQG control loop for the Reusable Launch Vehicle (RLV) aerodynamic descent.	33
Figure 3.4 – North-wind profile and turbulence characterization. (a) Mean northward wind $N(z)$. (b) Time series of the u (north) component . (c) Histograms/PDFs of u -gust (light and strong).	35
Figure 4.1 – Roll: Nominal open-loop step.	43
Figure 4.2 – Downrange: Nominal open-loop step.	44
Figure 4.3 – Crossrange: Nominal open-loop step.	44
Figure 4.4 – Roll: Augmented model, open-loop step.	45
Figure 4.5 – Downrange: Augmented model, open-loop step.	46
Figure 4.6 – Crossrange: Augmented model, open-loop step.	46
Figure 4.7 – Belém wind profile.	48
Figure 4.8 – Mixed northward wind profile (mean + gust).	49
Figure 4.9 – Step responses for the Downrange LQR controllers (five models).	52
Figure 4.10–Step responses for the scheduled Roll LQR controllers.	52
Figure 4.11–Step responses for Crossrange LQR controllers.	53
Figure 4.12–Singular values of T and S for Downrange.	53
Figure 4.13–Singular values of T and S for Roll.	54
Figure 4.14–Singular values of T and S for Crossrange.	54
Figure 4.15–Closed-loop Bode diagram with margins for Downrange.	55
Figure 4.16–Closed-loop Bode diagram with margins for roll.	56

Figure 4.17–Closed-loop Bode diagram with margins for Crossrange.	56
Figure 4.18–Step responses for the scheduled Downrange KF estimators.	58
Figure 4.19–Step responses for the scheduled roll KF estimators.	58
Figure 4.20–Step responses for the scheduled Crossrange KF estimators.	59
Figure 4.21–Singular values of T and S for Downrange.	59
Figure 4.22–Singular values of T and S for roll.	60
Figure 4.23–Singular values of T and S for Crossrange.	60
Figure 4.24–Closed-loop Bode diagram with margins for Downrange KF estimators.	61
Figure 4.25–Closed-loop Bode diagram with margins for roll KF estimators.	62
Figure 4.26–Closed-loop Bode diagram with margins for Crossrange KF estimators.	62
Figure 4.27–Step responses for scheduled Downrange LQG/LTR controllers.	63
Figure 4.28–Step responses for scheduled roll LQG/LTR controllers.	64
Figure 4.29–Step responses for scheduled Crossrange LQG/LTR controllers.	64
Figure 4.30–Singular values of T and S for Downrange.	65
Figure 4.31–Singular values of T and S for roll.	66
Figure 4.32–Singular values of T and S for Crossrange.	66
Figure 4.33–Closed-loop Bode diagram and margins for Downrange LQG/LTR.	67
Figure 4.34–Closed-loop Bode diagram and margins for roll LQG/LTR.	67
Figure 4.35–Closed-loop Bode diagram and margins for Crossrange LQG/LTR.	68
Figure 4.36–All-model overlay: reference tracking and control effort across the full trajectory.	70
Figure 4.37–All-model overlay: state evolution (pairs) over the same run.	71
Figure 4.38–Operating-point coverage along the descent: active model index by time (hence altitude), with synthesis points (black dots) and switching boundaries (light grey dashed lines).	72
Figure 4.39–Reference $r(t)$, measured output $y(t)$ (solid), and internal estimate $\hat{y}(t)$ (dashed) Nominal.	73
Figure 4.40–Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ Predictive.	74
Figure 4.41–Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ Weighted.	75
Figure 4.42–Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ under injected disturbances: Nominal.	77
Figure 4.43–Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ under injected disturbances: Predictive.	77
Figure 4.44–Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ under injected disturbances: Weighted.	78

List of tables

Table 2.1 – Common interpolation methods for gain scheduling	25
Table 4.1 – Controllability/Observability ranks by axis and model (augmented realization).	47
Table 4.2 – Models 1–5 and their defining flightcondition parameters.	50
Table 4.3 – LQR gain and phase margins at scheduled points.	57
Table 4.4 – KF estimators: gain and phase margins across scheduled models.	61
Table 4.5 – LQG/LTR gain and phase margins across scheduled models.	68
Table 4.6 – Optimal cost values J_{LQR} for control strategies	75
Table 4.7 – Performance without disturbance.	76
Table 4.8 – Optimal cost values J_{LQR} for control strategies under disturbance	78
Table 4.9 – Performance under disturbance.	79

List of Abbreviations

6-Dof	six-Degrees-of-Freedom
AGL	Above Ground Level
ARE	Algebraic Riccati Equation
CALLISTO	Cooperative Action Leading to Launcher Innovation in Stage Toss-back Operations
CFD	Computational Fluid Dynamics
CNES	French Space Agency
DARE	Discrete Algebraic Riccati Equation
DCA	Downrange–Crossrange–Altitude
DCM	Direction-Cosine Matrix
DLR	German Aerospace Center
GM	gain margin
GNC	Guidance, Navigation, and Control
JAXA	Japan Aerospace Exploration Agency
KF	Kalman Filter
LQG	Linear Quadratic Gaussian
LQR	Linear Quadratic Regulator
LTI	Linear Time-Invariant
LTR	Loop Transfer Recovery
MECO	Main Engine Cutoff
MEIG	Main Engine Ignition
NED	North–East–Down frame
PDF	Probability Density Function
PID	Proportional–Integral–Derivative
PM	phase margin
RLV	Reusable Launch Vehicle
UEN	Up–East–North
UFPA	Federal University of Pará
VTVL	Vertical Take-Off, Vertical Landing

List of Symbols

A	altitude position, m
$a_{i,j}$	(i, j) th element of the dynamic matrix, $[\text{DTM}(j) \cdot \text{DTM}(j)^{-1}]$
a^{acro}	aerodynamic acceleration vector, $\text{m} \cdot \text{s}^{-2}$
a^{fict}	fictitious acceleration vector, $\text{m} \cdot \text{s}^{-2}$
a^{grav}	gravity acceleration vector, $\text{m} \cdot \text{s}^{-2}$
$b_{i,j}$	(i, j) th element of the control matrix, $[\text{DTM}(j) \cdot \text{DTM}(j)^{-1}]$
C	crossrange position, m
C_I	torque aerodynamic coefficient around x
C_m	torque aerodynamic coefficient around y
C_n	torque aerodynamic coefficient around z
C_x	force aerodynamic coefficient along x
C_y	force aerodynamic coefficient along y
C_z	force aerodynamic coefficient along z
D	downrange position, m
ea	set of Euler angles with respect to downrange cross-range altitude, rad
F^{acro}	aerodynamic force vector, N
g_0	gravity acceleration at sea level, $\text{m} \cdot \text{s}^{-2}$
I	inertia matrix, $\text{kg} \cdot \text{m}^2$
I_{ij}	(i, j) th element of the inertia matrix, $\text{kg} \cdot \text{m}^2$
L_{ref}	reference length, m
M	Mach number
M_Δ	μ -analysis-related closed-loop system
M_{11}	left-upper submatrix of M
m	mass, kg
q	dynamic pressure, $\text{N} \cdot \text{m}^{-2}$
R_A^B	rotation matrix from A to B
R_{air}	specific air constant, $\text{J} \cdot \text{kg}^{-1} \cdot \text{K}^{-1}$
r	position vector, m
r_{LP}	position vector of landing pad with respect to the central body, m
S_{ref}	reference surface, m^2
T	atmospheric temperature, K
V	speed, $\text{m} \cdot \text{s}^{-1}$
V_A	altitude velocity, $\text{m} \cdot \text{s}^{-1}$
V_C	crossrange velocity, $\text{m} \cdot \text{s}^{-1}$
V_D	downrange velocity, $\text{m} \cdot \text{s}^{-1}$

V_{Cw}	crossrange wind velocity, $\text{m} \cdot \text{s}^{-1}$
V_{Dw}	downrange wind velocity, $\text{m} \cdot \text{s}^{-1}$
v	velocity vector, $\text{m} \cdot \text{s}^{-1}$
v^{air}	air-speed rocket velocity vector, $\text{m} \cdot \text{s}^{-1}$
v^{gs}	ground-speed rocket velocity vector, $\text{m} \cdot \text{s}^{-1}$
v^{wind}	wind velocity vector, $\text{m} \cdot \text{s}^{-1}$
W	generic weight matrix, $[\text{DTM}()]$ for input and $[\text{DTM}()^{-1}]$ for output
$(\cdot)_c$	commanded value, $[\text{DTM}()]$
$\dot{(\cdot)}$	first time derivative, $[\text{DTM}() \cdot \text{s}^{-1}]$
$(\cdot)_\oplus$	Earth-related value, $[\text{DTM}(\cdot)]$
α	angle of attack, rad
α_{gs}	ground-based tail angle of attack, rad
β	sideslip angle, rad
γ_{air}	air specific heat ratio
Δ	uncertainties diagonal matrix
δ_θ	pitch virtual deflection, rad
δ_ϕ	roll virtual deflection, rad
δ_ψ	yaw virtual deflection, rad
θ	pitch angle with respect to downrange crossrange altitude, rad
μ_\oplus	Earth's gravitational parameter, $\text{m}^3 \cdot \text{s}^{-2}$
ρ	atmospheric density, $\text{kg} \cdot \text{m}^{-3}$
ϕ	roll angle with respect to downrange crossrange altitude, rad
ψ	yaw angle with respect to downrange crossrange altitude, rad
ω	angular rate vector with respect to DCA in body, $\text{rad} \cdot \text{s}^{-1}$

Summary

1	Introduction	1
1.1	Context	1
1.2	Motivation	2
1.3	Objectives	4
1.3.1	Main Objective	4
1.3.2	Specific Objectives	4
1.4	Control Methods Background	5
1.5	Structure of the Document	8
2	Theoretical Foundation	9
2.1	Initial Considerations	9
2.2	Aerodynamic Modeling for Reusable Launch Vehicles	9
2.2.1	Vehicle and Mission Description	10
2.2.2	Reference Frames and Coordinates	11
2.3	Nonlinear 6-DoF Model	14
2.4	State Space Representation	15
2.4.1	Equations of Motion	15
2.4.2	Linearization and Trim Conditions	16
2.5	Discrete-Time LQR Design	16
2.5.1	Problem Formulation	16
2.5.2	Discrete Algebraic Riccati Equation	17
2.5.3	Interpretation and Properties	17
2.6	State Estimation: Kalman Filter	18
2.6.1	Design, tuning, and frequency insight	18
2.7	Linear Quadratic Gaussian (LQG) Design	18
2.7.1	Discrete-Time Plant Model	19
2.7.2	LQR Component	19
2.7.3	Kalman Filter Component	19
2.7.4	LQG Controller by Separation Principle	19
2.8	Loop Transfer Recovery	20
2.9	Robustness and Frequency-Domain Analysis	21
2.9.1	Robustness in control	21
2.9.2	Frequency-domain tools for robustness assessment	22
2.10	Scheduling	24
2.11	Summary	25
3	Methodology	27
3.1	Initial Considerations	27

3.2	Synthesis Point Identification for Scheduling	28
3.3	Model Augmentation and Pre-compensator Design	29
3.4	Proposed Controller Architecture	31
3.4.1	Atmosphere and Disturbances	33
3.5	Linearization and Discretization	36
3.5.1	Numerical Considerations	36
3.6	Summary	37
4	Results	39
4.1	Simulation Setup and Assumptions	39
4.1.1	Models and Implementation	39
4.1.2	Test Scenarios	50
4.1.3	Performance Indices	51
4.2	Robustness and Frequency-Domain Analysis	51
4.2.1	LQR Loop Shapes and Margins	51
4.2.2	Kalman Filter Loop Shapes and Stability Margins	57
4.2.3	LQG/LTR Loop Shapes and Stability Margins	63
4.2.4	LTR Turning Effect	68
4.3	Trajectory-Wide Cross-Regime Comparison	69
4.4	Scheduling and Implementation Aspects	72
4.4.1	Mapping and selection rule.	72
4.4.2	Reference-Tracking Refine Strategies under Time-Scheduled Control	73
4.4.3	Performance	75
4.5	Stress Testing and Sensitivity	76
4.6	Summary	79
5	Conclusions	80
5.1	Main Achievement	80
5.2	Future Works	82
	Appendix	85
	References	91

1 Introduction

Space exploration has fundamentally reshaped human civilization, not only technologically but also socially and philosophically. The modern era of space exploration began on October 4, 1957, when the Soviet Union successfully launched *Sputnik 1*, the first artificial satellite in history (OLIVEIRA, 2023). This groundbreaking achievement marked the inception of the Space Age and ignited an intense scientific and technological competition between major world powers, commonly referred to as the *Space Race*. The event captured global attention, stimulating unprecedented advancements in rocketry, orbital mechanics, and communication systems, and paving the way for human spaceflight, interplanetary probes, and the continuous expansion of our knowledge about the Universe.

Satellites now form the backbone of global infrastructure, enabling instant communication across continents, precise navigation for transportation, and real-time monitoring of Earth's climate and resources. Beyond practical applications, space missions have expanded our understanding of the cosmos, offering insights into planetary formation, the potential for extraterrestrial life, and humanity's place in the universe (NASA, 2024). However, the sustainability of space activities remains a critical challenge. The rise of Reusable Launch Vehicles (RLVs) promises to reduce costs and environmental footprints, yet achieving safe and reliable reusability demands breakthroughs akin to the Apollo program's ingenuity (OLIVEIRA; LAVAGNA, 2024). Looking ahead, space exploration continues to drive innovation, inspire international collaboration, and address existential questions proving that its utility extends far beyond technology, into the very fabric of human progress and aspiration.

1.1 Context

The reuse of rockets has emerged as one of the greatest revolutions in the history of space exploration, redefining the paradigms of space access. For decades, launch vehicles were entirely expendable, following a model inherited from the Apollo era and the Space Shuttle, which, although partially reusable, maintained prohibitive costs due to the complexity of the maintenance between missions (CRITICISM..., 2025). This reality began to change radically starting in 2015, when advances in engineering, flight control, and cost management enabled the practical adoption of reusable rockets, transforming what was once confined to science fiction into a viable commercial strategy.

SpaceX, a pioneer in this transition, demonstrated that recovering and reusing rocket stages could drastically reduce mission costs while increasing access to space. Its Falcon 9 became the symbol of this new era, with boosters landing vertically on ocean

platforms and being reused within months (SpaceX, 2015). Blue Origin achieved a similar breakthrough with the vertical landing of *New Shepard* and is currently developing the heavy-lift reusable vehicle *New Glenn* (ORIGIN, 2019). Rocket Lab introduced recovery attempts with *Electron* and announced its partially reusable *Neutron*, while Europe (ESA’s Themis program) and China (Long March 8R) also began investing in reusable technologies (DLR-CNES-JAXA, 2025).

Several collaborations are being established to support this paradigm shift. This is the case with Cooperative Action Leading to Launcher Innovation in Stage Toss-back Operations (CALLISTO), a pioneering international demonstrator for reusable launch vehicle technologies developed jointly by the German Aerospace Center (DLR), the French Space Agency (CNES), and the Japan Aerospace Exploration Agency (JAXA) (SAGLIANO et al., 2023). Similarly, the ESA-NASA Artemis Program brings together 30 nations to establish sustainable lunar exploration, while the US-ROK (Republic of Korea) partnership focuses on developing next-generation reusable rocket engines. The India-France Gaganyaan-French Space Agency (CNES) collaboration combines expertise in crewed spaceflight systems, and the UK-Australia Southern Launch initiative optimizes launch trajectories for reusable vehicles.

1.2 Motivation

While traditional expendable rockets operated under relatively simple dynamic regimes, with essentially open trajectories, reusable launch systems introduce multi-domain complexities that challenge conventional control paradigms. The need to precisely guide rocket stages through critical phases of atmospheric re-entry and vertical landing demands not only technological advances but also a fundamental reformulation of modeling and control approaches for these highly nonlinear and time-varying dynamic systems (HANSON, 2000; SAGLIANO et al., 2023; OLIVEIRA; LAVAGNA, 2024; KIM; LEE; PARK, 2023).

The core of this challenge lies in the multivariable, strongly coupled nature of reusable rocket dynamics. During return maneuvers, the vehicle encounters radically different aerodynamic regimes from hypersonic flight during re-entry to subsonic flight during final approach. Each of these phases exhibits distinct dynamic characteristics, with rapidly changing parameters, pronounced nonlinearities, and significant cross-axis coupling (OLIVEIRA; LAVAGNA, 2024).

The shift from expendable to reusable launch systems introduces a significant increase in dynamic complexity, particularly during descent and landing operations. Unlike traditional rockets, reusable vehicles must traverse a sequence of highly variable aerodynamic regimes, from *hypersonic re-entry* to *subsonic vertical landing*, each characterized by distinct nonlinear dynamics, multi-axis couplings, and uncertain, time-varying parameters.

The design of guidance and control systems for such missions must therefore reconcile high performance, robustness, and adaptability in the face of rapidly changing flight conditions and external disturbances.

In a flight environment characterized by aerodynamic uncertainties, flexible-body effects, and unpredictable atmospheric disturbances, the development of control systems for reusable rockets requires special attention to robustness and disturbance rejection, particularly during the critical aerodynamic descent phase (SAGLIANO et al., 2023). Controller robustness thus becomes imperative, requiring guaranteed stability and adequate performance even in the presence of substantial parametric variations, modeling uncertainties, and unforeseen external disturbances.

Classical control methods have historically provided the foundation for aerospace applications, with Proportional–Integral–Derivative (PID) controllers and LQR serving as standard baselines. PID controllers are widely appreciated for their simplicity and ease of implementation, which made them attractive for early flight systems and launch vehicles. However, their effectiveness degrades significantly in highly nonlinear and time-varying environments such as those encountered during re-entry. In these regimes, PIDs demand extensive manual tuning and offer limited robustness to modeling errors and atmospheric disturbances, as extensively discussed in classical texts on computer-controlled systems and aerospace simulation (ÅSTRÖM; WITTENMARK, 1997; STEVENS; LEWIS; JOHNSON, 2016).

The introduction of LQR brought a more systematic and mathematically optimal approach by minimizing a quadratic performance index through state feedback. While successful in many aerospace applications, including rocket engines and aircraft control, LQR remains sensitive to modeling inaccuracies, lacks inherent integral action, and struggles to maintain performance in the presence of large uncertainties (MUSGRAVE, 1991; STEVENS; LEWIS; JOHNSON, 2016). To overcome these limitations, robust methods such as H_∞ control have been investigated in recent decades. These techniques embed robustness margins in the design and thus provide improved resilience against uncertainties. Nevertheless, their application to RLV under re-entry conditions has shown that they may lead to conservative control efforts and require precise uncertainty modeling, a challenging task in practice (SAGLIANO et al., 2023; OLIVEIRA; LAVAGNA, 2024).

The limitations of classical control methods highlighted the need for frameworks that combine optimal performance with improved resilience to disturbances and modeling errors. Within this context, the LQG controller emerged as a powerful alternative by integrating the LQR for optimal state-feedback control with the KF for optimal state estimation. This architecture provided an elegant means to address sensor noise and model perturbations, and has been widely studied in aerospace applications (STEVENS; LEWIS; JOHNSON, 2016; MUSGRAVE, 1991). However, a well-known drawback of the classical

LQG approach is its lack of guaranteed robustness margins when exposed to unmodeled dynamics or significant parametric deviations. To mitigate this issue, the LTR procedure was introduced. By shaping the open-loop transfer function through appropriate design of the estimator dynamics, LTR restores desirable gain and phase margins, thereby improving robustness while preserving the optimality properties of the LQG controller (STEVENS; LEWIS; JOHNSON, 2016). This synthesis, known as LQG/LTR, has been recognized as a promising framework for advanced aerospace systems, including reusable launch vehicles, where robustness against uncertainties and high-frequency disturbances is essential.

By leveraging gain-scheduling strategies and potentially hybrid integration with predictive guidance schemes, LQG/LTR-based controllers can be adapted to track optimal re-entry and landing trajectories across multiple aerodynamic regimes. Such an approach not only provides a mathematically grounded solution but also aligns with operational constraints such as actuator limits, fuel optimisation, and rapid reusability turnaround requirements.

Accordingly, this dissertation is centered on the following core question:

How can a LQG/LTR controller be designed, adapted, and validated to ensure robust and optimal control of a reusable rocket during aerodynamic descent under nonlinear, uncertain, and highly variable flight conditions?

1.3 Objectives

1.3.1 Main Objective

The main objective of this research is to develop, adapt, and validate an advanced control strategy, based on the LQG/LTR framework, for the aerodynamic descent and vertical landing phase of a reusable rocket. Emphasis is placed on ensuring robust performance in the presence of strong nonlinearities, parametric uncertainties, sensor noise, and atmospheric disturbances.

1.3.2 Specific Objectives

- **To explore and evaluate advanced optimal control techniques**, with particular emphasis on LQG and LQG/LTR architectures, as viable alternatives for the aerodynamic descent phase of reusable launch vehicles.
- **To model and analyze the nonlinear flight dynamics** of reusable rockets during re-entry and descent, identifying critical dynamic regimes and proposing linearized state-space models suitable for Linear Time-Invariant (LTI)-based control design.

- **To study the technical and theoretical challenges** inherent in controlling highly nonlinear and multi-regime flight trajectories under uncertainty, including actuator limitations, time-varying dynamics, and environmental disturbances.
- **To investigate the robustness characteristics of optimal controllers** using frequency domain analysis, including the sensitivity and complementary sensitivity functions, with special attention to non-minimum phase behaviors and the impact of modeling errors.
- **To propose a systematic approach for tuning the covariance matrices Q and R** in the Kalman filter, aiming to enhance robustness properties of LQG/LTR controllers under realistic noise and uncertainty profiles.
- **To design simulation frameworks and scenarios** (*six – Degrees – of – Freedom*(6 – *DoF*) or hybrid nonlinear/linear models) for validating the control strategies against representative perturbations and trajectory constraints in the descent phase.
- **To contribute theoretical and applied material** that supports the academic and scientific community in the study of reusable launch vehicle control, providing a foundation for further research and experimental implementation.

1.4 Control Methods Background

Understanding the historical evolution of control methodologies is essential for appreciating the architectures currently employed in aerospace Guidance, Navigation, and Control (GNC). The trajectory from classical to optimal and finally to robust control reflects both advances in theory and the lessons learned from decades of practical implementation in increasingly demanding missions. Each paradigm addressed shortcomings of its predecessors, while introducing new challenges of its own. The following review presents a continuous narrative of classical, optimal, and robust control approaches, with emphasis on their relevance to launch vehicles and reusable systems.

The development of control methods in aerospace has followed a trajectory shaped by the competing demands of simplicity, reliability, and performance under uncertainty. Early architectures relied almost exclusively on classical approaches, particularly PID controllers and frequency-domain design methods such as Bode plots and root locus analysis. These tools offered engineers intuitive graphical techniques to tune feedback laws and guarantee stability margins. They became the backbone of autopilots, missile stabilization systems, and the first space launchers. A key advantage of classical controllers was their transparency: parameters had a clear physical interpretation, and certifying

authorities could easily verify their behavior. The Space Shuttle relied heavily on such gain-scheduled controllers, with linear feedback loops designed around multiple trim conditions for ascent, on-orbit maneuvering, entry, and landing. While these controllers provided robustness in practice, they required extensive empirical adjustment, and any change in mass properties, aerodynamics, or mission profile imposed costly redesign and lengthy verification campaigns (STEVENS; LEWIS; JOHNSON, 2016; HANSON, 2000). The very success of classical control in enabling flight also revealed its limits, particularly for reusable launch vehicles where adaptability and reduced operations cost were paramount.

The limitations of tuning-intensive design motivated the adoption of optimal control frameworks during the 1960s and 1970s. The LQR provided a systematic state-space approach that minimized a quadratic cost, balancing tracking accuracy against control effort. Its solution, obtained from the Algebraic Riccati Equation (ARE), gave explicit guarantees of stability and optimality under nominal conditions. Combined with Kalman filtering for state estimation in noisy environments, this led to the LQG controller, one of the most influential methods in aerospace. LQG allowed multivariable control designs that could coordinate multiple actuators and explicitly handle process and sensor noise. These theoretical advances quickly found application in aeroengine regulation, flight control augmentation, and launcher attitude control (MUSGRAVE, 1991). However, Doyle's famous observation that there are no guaranteed stability margins for LQG revealed a vulnerability: optimal controllers for the nominal model could fail dramatically in the presence of unmodeled dynamics. LTR was introduced to address this issue, shaping the open-loop response to resemble that of an LQR and thereby recovering robustness margins. Nevertheless, the recognition that optimal controllers lacked intrinsic guarantees of robustness set the stage for a new generation of methods that could explicitly incorporate uncertainty (ÅSTRÖM; WITTENMARK, 1997).

In practice, aerospace engineers continued to employ both classical and optimal methods. Classical control remained attractive for single-input single-output loops, where its simplicity and ease of certification were decisive. Optimal methods, on the other hand, provided valuable tools for multivariable systems and for coordinating actuators such as thrust vectoring, aerodynamic surfaces, and reaction jets. Studies on advanced reusable launchers, including NASA's X-33 and VentureStar demonstrators, highlighted the tradeoffs. Flight control analysis showed that gain-scheduled controllers could maintain stability across wide flight envelopes but required significant empirical tuning, while LQG/LTR approaches offered systematic synthesis but demanded augmentation to achieve practical robustness (CALHOUN, 2000; HANSON, 2000). This duality reinforced the view that no single method was sufficient and that future architectures would require hybrid philosophies integrating classical transparency, optimal efficiency, and robust guarantees.

Robust control emerged in the late twentieth century to close this gap. Its central

objective was to guarantee stability and performance under model uncertainty and external disturbances, conditions that are unavoidable in aerospace applications. Unlike classical margins, which are indirect indicators of robustness, and unlike LQG designs, which optimize only the nominal case, robust synthesis incorporated uncertainty explicitly into the design problem. The H_∞ framework became the cornerstone of this approach. By formulating control as the minimization of the worst-case gain from disturbances to outputs, H_∞ ensured that system responses remained bounded even under adverse conditions. The introduction of weighting functions allowed designers to specify desired performance in particular frequency ranges, making the methodology highly flexible. These features made H_∞ appealing for flight control, vibration suppression, and launcher guidance problems (ÅSTRÖM; WITTENMARK, 1997).

However, practical application of H_∞ encountered challenges. The controllers generated were often of very high order, sometimes as large as the augmented plant itself, complicating implementation and certification. This limitation gave rise to structured H_∞ synthesis, in which the controller is restricted to a predefined form such as a PID, gain-scheduled law, or low-order dynamic compensator. By constraining the architecture, designers could retain the transparency needed for flight certification while inheriting the robustness guarantees of H_∞ . Structured methods proved especially valuable in launcher GNC, where certifiable low-order controllers are required but robustness to aerodynamic and structural uncertainties is indispensable. Recent applications confirm this trend: Sagliano et al. (2023) demonstrated unified structured H_∞ controllers for CALLISTO, a reusable launch demonstrator developed by German Aerospace Center (DLR), Japan Aerospace Exploration Agency (JAXA), and CNES, achieving resilience against uncertainties while maintaining pinpoint landing accuracy.

The evolution of robust methods also parallels advances in computational guidance. The development of convex optimization and successive convexification enabled realtime generation of descent trajectories that respect constraints on thrust, aerodynamics, and landing accuracy. Coupling these guidance algorithms with robust structured H_∞ controllers has produced architectures capable of ensuring both trajectory feasibility and closed-loop stability under uncertainty. De Oliveira e Lavagna (2024), for instance, showed that such integration allows reusable launchers to perform precise landings even under significant perturbations, outperforming classical gain-scheduling and baseline optimal approaches. This integration highlights the direction of current research: robustness is no longer treated in isolation but as a component of a unified GNC framework that also encompasses guidance and trajectory optimization.

Taken together, the trajectory from classical to optimal to robust control reflects the maturation of aerospace engineering. Classical methods delivered the first reliable flight controllers, intuitive to design and certify but costly to adapt. Optimal control introduced

mathematical rigor and multivariable capability, but its vulnerability to uncertainty limited its standalone use. Robust control, particularly H_∞ and its structured variants, provided a synthesis that could reconcile performance with guaranteed margins. In modern practice, these approaches are not mutually exclusive but complementary: classical laws often provide baseline loops, optimal methods coordinate multivariable dynamics, and robust synthesis ensures resilience. The convergence of these traditions has produced today's GNC architectures for reusable launch vehicles, where adaptability, safety, and cost efficiency must coexist.

1.5 Structure of the Document

This dissertation is organized into five chapters, structured to guide the reader from the theoretical foundations to the practical implementation and analysis of the proposed control strategies.

Chapter 1 introduces the general context of the research, presenting the scientific and technological motivation, provides a comprehensive literature review, examining previous research efforts related to the control of reusable launch vehicles, with particular focus on classical and modern control approaches, as well as the main and specific objectives that define the scope of this work.

Chapter 2 covers the theoretical foundations necessary for the development of this study. It includes the dynamic modeling of a fin-stabilized rocket, the design of LQR controllers, Kalman filter formulation, LQG controller synthesis, the Loop Transfer Recovery technique, and relevant robustness criteria based on frequency-domain performance analysis.

Chapter 3 details the methodology adopted for controller implementation and simulation. This includes the linearization and discretization procedures, definition of test scenarios, selection of performance indices, and the numerical tools used throughout the development.

Chapter 4 presents and analyzes the simulation results, including graphical representations of system responses, comparative performance between control architectures, and discussion based on key metrics such as stability, accuracy, and robustness.

Finally, **Chapter 5** summarizes the main contributions of the research, highlighting the theoretical and practical findings, and outlines potential directions for future work, including the experimental validation of the proposed strategies.

2 Theoretical Foundation

This chapter presents the theoretical concepts and mathematical tools forming the foundation for control design of RLVs during aerodynamic flight phases. This phase is characterized by strong nonlinearities, multi-axis coupling, and time-varying parameters. The following sections provide the aerodynamic modeling background, state-space formulation, and the principles of optimal and robust control strategies.

2.1 Initial Considerations

The descent phase of an RLV is highly challenging due to the continuous change in aerodynamic regime: from supersonic at the upper atmosphere, through transonic near max dynamic pressure, down to subsonic during final approach (HANSON, 2000). These transitions significantly alter aerodynamic forces, stability derivatives, and actuator effectiveness. In addition, atmospheric disturbances such as wind shear, turbulence, and density variations must be managed in real time.

From a control perspective, these conditions require:

- Models that capture the essential dynamics in each regime;
- Control strategies that ensure robustness to large parameter variations;
- Estimation schemes that remain accurate despite sensor noise and model mismatch.

The aim of this chapter is to present the theoretical tools enabling such capabilities.

2.2 Aerodynamic Modeling for Reusable Launch Vehicles

Recent research has emphasized the need for aerodynamic models of RLV that can accurately capture the strong nonlinearities and parameter variations occurring throughout the atmospheric descent. Sagliano et al. (2023) developed a 6-DoF aerodynamic model for the CALLISTO demonstrator, incorporating variations in aerodynamic coefficients with Mach number, angle of attack, and control surface deflections. The model integrates data from Computational Fluid Dynamics (CFD) analyses and wind-tunnel tests, providing high-fidelity aerodynamic force and moment predictions across the supersonic, transonic, and subsonic regimes. Importantly, the formulation includes the coupling between longitudinal and lateral dynamics, enabling realistic simulation of multi-axis control challenges. This modeling framework also accounts for environmental disturbances such as wind gusts

and atmospheric density variations, making it suitable for evaluating robust control strategies under realistic flight conditions. These characteristics are critical for designing and validating controllers that can maintain trajectory accuracy and stability despite rapidly changing aerodynamic conditions during descent.

The aerodynamic modeling of the RLV considered in this work is based on the 6-DoF nonlinear equations of motion of the CALLISTO demonstrator, as presented by Sagliano et al. (2023). This model captures the vehicle dynamics across multiple flight regimes—from supersonic re-entry to subsonic vertical approach, and provides the foundation for the synthesis and evaluation of robust controllers.

This section presents a control-oriented aerodynamic model for RLVs covering supersonic, transonic, and subsonic regimes. The formulation supports robust and optimal controller synthesis by exposing state-space structure, uncertainty descriptions, and scheduling variables.

2.2.1 Vehicle and Mission Description

The CALLISTO vehicle is a Vertical Take-Off, Vertical Landing (VTVL) booster demonstrator equipped with four independently actuated aerodynamic fins for atmospheric control. It is a small-scale experimental vehicle, about 13 m tall and 1.1 m in diameter, designed for suborbital flights and incorporates throttleable engines, guidance and control algorithms, and landing gear to perform multiple flights with rapid refurbishment.

Thanks to a set of four aerodynamic fins, control over three axes, Z, Y and X, is obtained during non-propelled phases. These four physical fins can be mapped onto an equivalent set of three virtual fins δ_ϕ , δ_θ , and δ_ψ , which are used to express the aerodynamic coefficients with respect to the three possible rotations and, therefore, are actively used to formulate the equations of motion. This mapping is given by (2.1) preserving the decoupling between axes in first approximation (SAGLIANO et al., 2023).

$$\begin{bmatrix} \delta_\phi \\ \delta_\theta \\ \delta_\psi \end{bmatrix} = \begin{bmatrix} -1/4 & -1/4 & -1/4 & -1/4 \\ 1/2 & 0 & -1/2 & 0 \\ 0 & 1/2 & 0 & -1/2 \end{bmatrix} \begin{bmatrix} \delta_1 \\ \delta_2 \\ \delta_3 \\ \delta_4 \end{bmatrix} \quad (2.1)$$

At the end of the ascent phase the vehicle begins its boostback maneuver to invert the horizontal direction of the velocity vector. After this maneuver is completed a main engine cutoff (MECO) command is issued, and the vehicle starts its unpowered descent. The actively controlled aerodynamic descent begins as soon as the dynamic pressure is strong enough. In this phase, the attitude control is ensured by the fins that will steer the vehicle such that a given aerodynamic force, needed to track the trajectory, is generated.

The aerodynamic control phase is concluded at a specific main engine ignition (MEIG) altitude, where the engine is then re-activated and the powered descent and landing phase begins. The exclusive focus in this work, is on the tracking of the comeback trajectory during aerodynamic descent phase. An illustration of the mission and its different phases of flight are shown in Figure 2.1. Sagliano et al. (2023) state that the aerodynamic descent spans from 25 km down to 5 km in altitude, with a total duration of around 60 seconds.

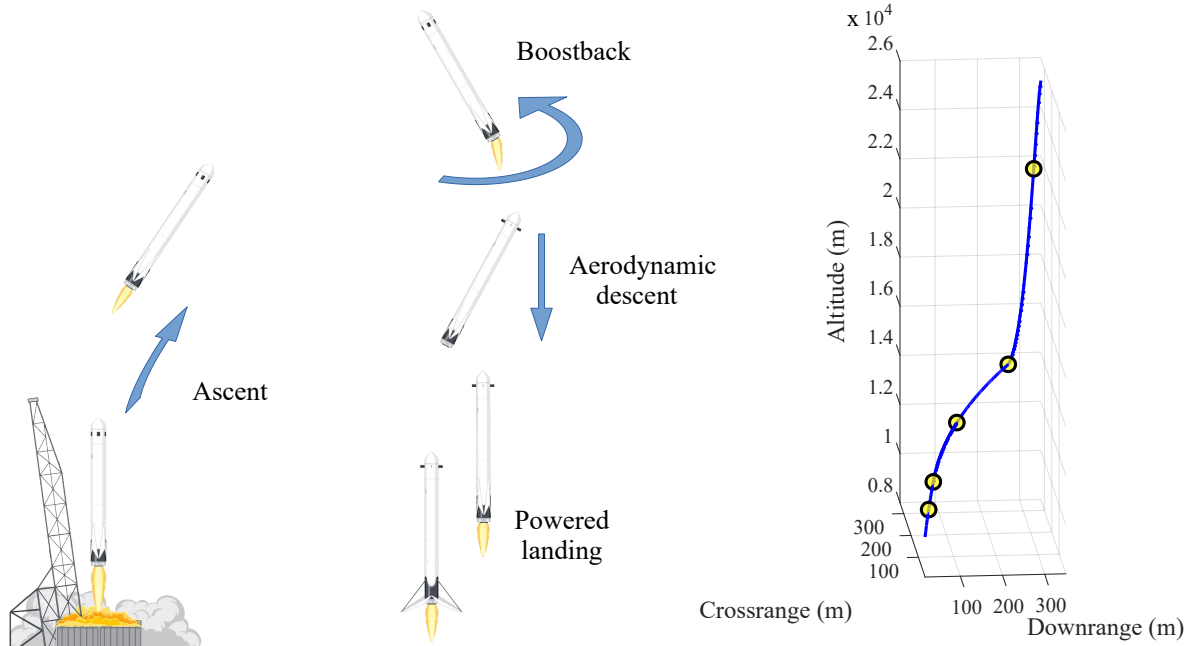


Figure 2.1 – Mission illustration in left and 3D plotting of a suggested trajectory for back to launch site scenario by right.

Source: Author

2.2.2 Reference Frames and Coordinates

The accurate definition of reference frames is a prerequisite for modeling the 6-DoF dynamics of a RLV. Following the framework adopted in recent studies on aerodynamic descent and precision landing (SAGLIANO et al., 2023; OLIVEIRA; LAVAGNA, 2024), two principals reference frames are employed to describe the vehicle dynamics during descent. The first is a *local inertial frame* \mathcal{I} , centered at the landing site, whose axes follow an Up–East–North (UEN) convention: the $x_{\mathcal{I}}$ -axis points upward, the $y_{\mathcal{I}}$ -axis eastward, and the $z_{\mathcal{I}}$ -axis northward. For the duration of the descent, this frame is treated as inertial, providing a fixed reference for trajectory analysis and landing accuracy. It is not completely aligned with the first frame, but is obtained by rotating the UEN frame about the local vertical by an angle $\mathcal{X}_{\mathcal{I}}$ to align the D -axis with the trajectory, as illustrated in Figure 2.2.

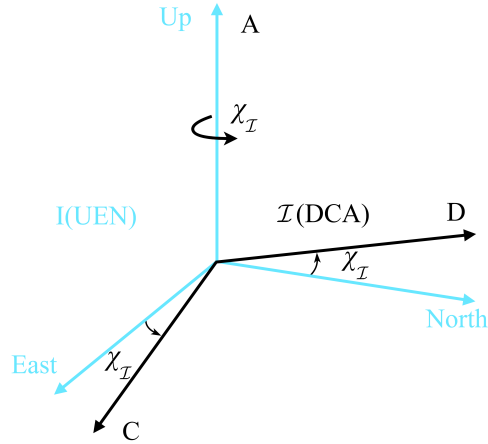


Figure 2.2 – Inertial UEN and DCA references frame emphasizing χ_I rotation to decoupling.

Source: Author

The *body-fixed* frame \mathcal{B} is attached to the vehicles center of mass as shown in Figure 2.3. Its axes are defined as follows: X_B points from the CoM toward the *nose* (up), Z_B lies in the vehicles longitudinal symmetry plane, and Y_B completes a right-handed triad by definition (2.2)

$$Y_B = Z_B \times X_B \quad (2.2)$$

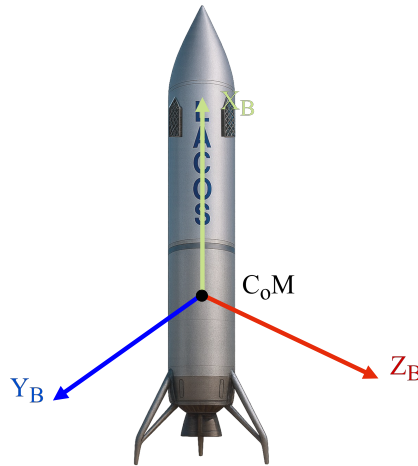


Figure 2.3 – Body references frame.

Source: Author

The orientation of \mathcal{B} with respect to the inertial frame \mathcal{I} is classically parameterized by ZYX Euler angles (ϕ, θ, ψ) , interpreted here as roll about X_B , pitch about Y_B , and yaw about Z_B . To avoid singularities, attitude kinematics are propagated with the unit quaternion $q_{\mathcal{I}\mathcal{B}}(t)$ and the associated direction-cosine matrix (DCM) $R_B^{\mathcal{I}}(t)$ described in (2.3). Angular velocities $\boldsymbol{\omega} = [p \ q \ r]^T$ are expressed in the body-fixed frame, providing a

consistent six-degree-of-freedom formulation.

$$R_{\mathcal{B}}^{\mathcal{I}} = \begin{bmatrix} \cos\theta\cos\psi & \cos\theta\sin\psi & -\sin\theta \\ -\cos\phi\sin\psi + \sin\phi\sin\theta\cos\psi & \cos\phi\cos\psi + \sin\phi\sin\theta\sin\psi & \sin\phi\cos\theta \\ \sin\phi\sin\psi + \cos\phi\sin\theta\cos\psi & -\sin\phi\cos\psi + \cos\phi\sin\theta\sin\psi & \cos\phi\cos\theta \end{bmatrix} \quad (2.3)$$

Remark. In descent phases, a practical alignment with the DCA frame often yields $X_{\mathcal{B}} \approx D$, $Y_{\mathcal{B}} \approx C$, and $Z_{\mathcal{B}} \approx -A$, although the above definitions do not depend on that choice. But the actual context, vertical descent phase, the alignment follows: $X_{\mathcal{B}} \approx A$, $Y_{\mathcal{B}} \approx C$, and $Z_{\mathcal{B}} \approx D$.




Figure 2.4 – 3-2-1 Rotation sequence to move from Downrange-Crossrange-Altitude to body axes.

Source: Author

The wind-path frame \mathcal{W} is an *auxiliary* axes set aligned with the air-relative velocity \mathbf{V}_{rel} . It is used to define the aerodynamic angles (α, β) and to express aerodynamic forces and moments as functions of these angles as shown Figure 2.5, while the equations of motion are typically propagated in body or navigation frames (STEVENS; LEWIS; JOHNSON, 2016; SAGLIANO et al., 2023).




Figure 2.5 – Representation of body reference frame, air-path reference frame, angle of attack, and sideslip angle.

Source: Author

In this framework, the aerodynamic force and moment coefficients $C_L, C_D, C_C, C_m, C_\ell, C_n$ are most naturally expressed as functions of the variables, angle of attack α , sideslip β , Mach number M , and altitude h . The state variables, α and β , coincide with the ISO-1151 geometric definitions (projection of \mathbf{V}_{rel} onto the $X_B - Z_B$ plane for α , and deviation from that plane for β) (STEVENS; LEWIS; JOHNSON, 2016).

2.3 Nonlinear 6-DoF Model

Building on the modeling assumptions and the adopted reference frames, the 6-DoF dynamics of RLV are expressed by the system of equations (2.6). This system will subsequently be linearized under suitable simplifying assumptions to derive local linear models at selected synthesis points, which are used for simulation and control-law design.

For analysis we use the Euler mapping from body rates $\boldsymbol{\omega} = [p q r]^\top$ to angle rates,

$$\dot{\boldsymbol{\Theta}} = \underbrace{\begin{bmatrix} 1 & \sin \phi \tan \theta & \cos \phi \tan \theta \\ 0 & \cos \phi & -\sin \phi \\ 0 & \sin \phi / \cos \theta & \cos \phi / \cos \theta \end{bmatrix}}_{\mathbf{T}(\phi, \theta)} \boldsymbol{\omega} \quad (2.4)$$

while for time integration we favour quaternion kinematics to avoid singularities:

$$\dot{q} = \frac{1}{2} \Omega(\boldsymbol{\omega}) q, \quad \Omega(\boldsymbol{\omega}) = \begin{bmatrix} 0 & -p & -q & -r \\ p & 0 & r & -q \\ q & -r & 0 & p \\ r & q & -p & 0 \end{bmatrix} \quad (2.5)$$

The DCM $R_{\mathcal{I}}^{\mathcal{B}}(q) \in \text{SO}(3)$ maps vectors from \mathcal{I} to \mathcal{B} .

The rigid-body translational and rotational dynamics of the unpowered vehicle are

$$\begin{aligned} m \dot{\mathbf{v}} &= \mathbf{F}_g + \mathbf{F}_a + \mathbf{F}_p \\ \mathbf{J} \dot{\boldsymbol{\Theta}} + \boldsymbol{\omega} \times (\mathbf{J} \boldsymbol{\omega}) &= \mathbf{M}_a + \mathbf{M}_p \end{aligned} \quad (2.6)$$

with mass m , inertia tensor \mathbf{J} , gravity $\mathbf{F}_g = R_{\mathcal{I}}^{\mathcal{B}} [0 \ 0 \ -mg]^\top$, aerodynamic forces/moments ($\mathbf{F}_a, \mathbf{M}_a$), and possible propulsive contributions ($\mathbf{F}_p, \mathbf{M}_p$) (e.g. thrust-vector control during transitional phases). Position and velocity in \mathcal{I} satisfy $\dot{\mathbf{r}} = \mathbf{v}$; the air-relative velocity is $\mathbf{v}_{\text{air}} = \mathbf{v} - \mathbf{w}$, with wind \mathbf{w} expressed in \mathcal{I} and rotated as needed.

Aerodynamic loads use the coefficient form with dynamic pressure $q_{D_{\text{dyn}}} = \frac{1}{2} \rho V^2$, reference area S_{Ref} and length l_{Ref} (all in body axes):

$$\mathbf{F}_a = q_{D_{\text{dyn}}} S_{\text{Ref}} \begin{bmatrix} C_X \\ C_Y \\ C_Z \end{bmatrix}, \quad \mathbf{M}_a = q_{D_{\text{dyn}}} S_{\text{Ref}} l_{\text{Ref}} \begin{bmatrix} C_l \\ C_m \\ C_n \end{bmatrix} \quad (2.7)$$

$$C_{\star}(\alpha, \beta, M, R_e, \boldsymbol{\delta}) = C_{\star}^{(0)}(\alpha, \beta, M, R_e) + \sum_i \Delta C_{\star}^{(i)}(\alpha, \beta, M, R_e) \delta_i \quad (2.8)$$

with $\star \in \{X, Y, Z, m, \ell, n\}$ and $i \in \{\psi, \theta, \phi\}$.

Stability and damping enter through derivatives such as $C_{m_\alpha}, C_{n_\beta}, C_{l_p}, C_{m_q}, C_{n_r}$. Coefficients are tabulated over (α, β, M, R_e) and interpolated during simulation; control effectiveness is captured by $\partial C_{\star} / \partial \delta_i$.

2.4 State Space Representation

2.4.1 Equations of Motion

The 6-DOF nonlinear dynamics of an RLV can be expressed as:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}, t) \quad (2.9)$$

where \mathbf{x} contains position, velocity, attitude, and angular rates, and \mathbf{u} represents control surface deflections or thrust vectoring commands.

Following the modeling approach adopted for CALLISTO, we arrive at the formulation below.

$$\begin{bmatrix} \dot{\phi} \\ \dot{\omega}_x \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ a_{\omega_x, \phi} & a_{\omega_x, \omega_x} \end{bmatrix} \begin{bmatrix} \phi \\ \omega_x \end{bmatrix} + \begin{bmatrix} 0 \\ b_{\omega_x} \end{bmatrix} \delta_\phi \quad (2.10)$$

$$\begin{bmatrix} \dot{D} \\ \dot{v}_D \\ \dot{\theta} \\ \dot{\omega}_y \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ a_{v_D, D} & a_{v_D, v_D} & a_{v_D, \theta} & 0 \\ 0 & 0 & 0 & -1 \\ a_{\omega_y, D} & a_{\omega_y, v_D} & a_{\omega_y, \theta} & 0 \end{bmatrix} \begin{bmatrix} D \\ v_D \\ \theta \\ \omega_y \end{bmatrix} + \begin{bmatrix} 0 \\ b_{v_D, \delta_\theta} \\ 0 \\ b_{\omega_y, \delta_\theta} \end{bmatrix} \delta_\theta \quad (2.11)$$

$$\begin{bmatrix} \dot{C} \\ \dot{v}_C \\ \dot{\psi} \\ \dot{\omega}_z \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ a_{v_C, C} & a_{v_C, v_C} & a_{v_C, \psi} & 0 \\ 0 & 0 & 0 & -1 \\ a_{\omega_z, C} & a_{\omega_z, v_C} & a_{\omega_z, \psi} & 0 \end{bmatrix} \begin{bmatrix} C \\ v_C \\ \psi \\ \omega_z \end{bmatrix} + \begin{bmatrix} 0 \\ b_{v_C, \delta_\psi} \\ 0 \\ b_{\omega_z, \delta_\psi} \end{bmatrix} \delta_\psi \quad (2.12)$$

Since the aerodynamic descent is performed without real altitude control (vertical motion), but only attitude control (Euler angles) and the two other relative positions (downrange and crossrange), the equations above allow for separate analysis of pitch, roll, and yaw. Nevertheless, it is critical to underscore that the coupled dynamics between rotational and translational motion persist, particularly in pitch (governing longitudinal motion, i.e., downrange) and yaw (dictating lateral motion, i.e., crossrange). (2.10) enables roll analysis, while (2.11) facilitates pitch analysis. Yaw analysis is performed by (2.12). The expanded forms of all terms in these equations are detailed in Appendix A.

2.4.2 Linearization and Trim Conditions

For control design we linearize about operating points along a nominal descent. A trim $(\bar{\mathbf{x}}, \bar{\mathbf{u}})$ solves $f(\bar{\mathbf{x}}, \bar{\mathbf{u}}) = \mathbf{0}$ with $\bar{\mathbf{y}} = h(\bar{\mathbf{x}}, \bar{\mathbf{u}})$. Linearization yields:

$$\delta \dot{\mathbf{x}} = A(\rho) \delta \mathbf{x} + B(\rho) \delta \mathbf{u}, \quad \delta \mathbf{y} = C(\rho) \delta \mathbf{x} + D(\rho) \delta \mathbf{u} \quad (2.13)$$

$$A = \left. \frac{\partial f}{\partial \mathbf{x}} \right|_{(\bar{\mathbf{x}}, \bar{\mathbf{u}})}, \quad B = \left. \frac{\partial f}{\partial \mathbf{u}} \right|_{(\bar{\mathbf{x}}, \bar{\mathbf{u}})}, \quad C = \left. \frac{\partial h}{\partial \mathbf{x}} \right|_{(\bar{\mathbf{x}}, \bar{\mathbf{u}})}, \quad D = \left. \frac{\partial h}{\partial \mathbf{u}} \right|_{(\bar{\mathbf{x}}, \bar{\mathbf{u}})} \quad (2.14)$$

The scheduling¹ vector $\rho = \{M, \alpha, q, h\}$ encodes the operating condition; a small set of LTI models $\{A, B, C, D\}$ is extracted at representative points (e.g. changes in M and dynamic pressure q) and used for gain scheduling. Inter-model interpolation (e.g. barycentric over (M, q)) provides smooth parameter dependence. A practical rule is to place grid points so the rate-of-change of ρ remains well below the closed-loop bandwidth, limiting scheduling transients.

2.5 Discrete-Time LQR Design

The LQR is one of the most fundamental tools in optimal control, providing a systematic framework for designing feedback laws that balance performance and control effort. In aerospace engineering, LQR design has been extensively used for flight control, propulsion systems, and more recently for RLV guidance and descent problems. While the continuous-time formulation is well known, the discrete-time version DLQR is particularly relevant for digital implementation, where control signals are generated by on-board computers operating at fixed sampling intervals (ÅSTRÖM; WITTENMARK, 1997; STEVENS; LEWIS; JOHNSON, 2016).

2.5.1 Problem Formulation

Consider a linear time-invariant system in discrete time:

$$x_{k+1} = Ax_k + Bu_k \quad (2.15)$$

where $x_k \in \mathbb{R}^n$ is the state vector, $u_k \in \mathbb{R}^m$ is the control input, and A, B are constant matrices. The LQR problem seeks a state feedback control law

$$u_k = -Kx_k \quad (2.16)$$

that minimizes the infinite-horizon quadratic cost

$$J = \sum_{k=0}^{\infty} (x_k^\top Q x_k + u_k^\top R u_k) \quad (2.17)$$

¹ Notion described in the last section of this chapter.

with symmetric weighting matrices $Q \succeq 0$ and $R \succ 0$. The matrix Q penalizes deviations from the desired equilibrium, while R penalizes control effort, allowing designers to tune the aggressiveness of the controller.

2.5.2 Discrete Algebraic Riccati Equation

Finite-horizon backward Riccati recursion.

For the discrete LQR problem with quadratic cost J , the *backward Riccati recursion* reads

$$K_k = \left(R + B^\top P_{k-1} B \right)^{-1} B^\top P_{k-1} A \quad (2.18)$$

$$P_k = Q + A^\top P_{k-1} A - A^\top P_{k-1} B K_k \quad (2.19)$$

With the time-varying optimal gain K_k from (2.18), the closed-loop evolves as

$$x_{k+1} = \left(A - B K_k \right) x_k \quad (2.20)$$

and is stabilized over the horizon under the standard assumptions (stabilizability of (A, B) or at least detectability of $(Q^{1/2}, A)$).

2.5.3 Interpretation and Properties

The discrete LQR design provides several desirable properties. It guarantees optimality with respect to the quadratic cost, explicitly balancing performance and effort. The closed-loop poles are placed within the unit circle, ensuring discrete-time stability. Moreover, the structure of the Discrete Algebraic Riccati Equation (DARE) highlights the trade-off: a large Q produces a more aggressive controller with faster response but higher input demands, whereas a larger R yields smoother control at the expense of tracking precision. From a frequency-domain standpoint, classical LQR synthesis typically yields satisfactory gain and phase margins; however, when the state feedback is paired with a KF to form LQG, no explicit robustness guarantees follow. This limitation motivated extensions such as the LQG/LTR methodology, which recovers desirable open-loop transfer characteristics and robustness margins (MUSGRAVE, 1991; STEVENS; LEWIS; JOHNSON, 2016).

The discrete-time LQR provides a powerful, elegant, and implementable methodology for designing optimal state feedback controllers in aerospace applications. Its foundation in the Riccati equation, together with its guaranteed stability and tunable performance, make it a reference tool for both academic research and practical implementation in digital flight control systems. Although robustness limitations persist, DLQR continues to form the basis for advanced extensions such as LQG/LTR ensuring its enduring relevance in the design of guidance and control systems.

2.6 State Estimation: Kalman Filter

State estimation is central to flight control and guidance for aerospace systems. In practice, measurements are noisy and the full state vector is rarely available; the *discrete-time* KF provides the minimum-variance linear estimate of the state under Gaussian noise assumptions, and constitutes the estimation block of the LQG architecture via the separation principle (ÅSTRÖM; WITTENMARK, 1997; STEVENS; LEWIS; JOHNSON, 2016). In RLV applications, the Kalman filter fuses heterogeneous sensors (e.g., IMU, GPS, altimeters, aerodynamic sensors) to deliver consistent state information during demanding phases such as aerodynamic descent and precise landing.

We consider the linear, time-invariant sampled model

$$\begin{aligned}x_{k+1} &= Ax_k + Bu_k + w_k \\y_k &= Cx_k + v_k\end{aligned}\tag{2.21}$$

with zero-mean, white, mutually independent Gaussian noises $w_k \sim \mathcal{N}(0, Q_w)$ and $v_k \sim \mathcal{N}(0, R_v)$, and an initial prior $\hat{x}_{0|0}$ with covariance $P_{0|0} \succeq 0$. Matrices (A, C) are assumed detectable so that estimation error converges.

If (A, C) is detectable and $(A, Q_w^{1/2})$ is stabilizable, the covariance converges to the stabilizing solution $P \succeq 0$ of the discrete Riccati equation

$$P_{k+1} = A_d P_k A_d^\top + Q_w - A_d P_k C_d^\top (C_d P_k C_d^\top + R_v)^{-1} C_d P_k A_d^\top\tag{2.22}$$

yielding a constant *steady-state* Kalman gain

$$L = A_d P C_d^\top (C_d P C_d^\top + R_v)^{-1}\tag{2.23}$$

2.6.1 Design, tuning, and frequency insight

The designer tunes Q_w and R_v to reflect process and measurement uncertainty. Increasing Q_w (for given R_v) enlarges L and the estimator bandwidth, tracking fast state changes at the expense of higher noise sensitivity; increasing R_v does the opposite. From a frequency-domain viewpoint, a high-bandwidth estimator improves low-frequency disturbance reconstruction but can amplify high-frequency sensor noise and unmodeled dynamics (STEVENS; LEWIS; JOHNSON, 2016). In LQG designs, the separation principle allows the KF (from (A, C, Q_w, R_v)) and the LQR (from (A, B, Q, R)) to be designed independently while guaranteeing optimality for the combined stochastic problem (ÅSTRÖM; WITTENMARK, 1997; STEVENS; LEWIS; JOHNSON, 2016).

2.7 Linear Quadratic Gaussian (LQG) Design

The LQG controller represents a cornerstone of modern optimal control. It combines the optimal regulation of the LQR with the optimal state estimation of the Kalman

filter, resulting in a coherent controller suitable for systems with noisy measurements and stochastic disturbances. The discrete-time formulation is particularly relevant for aerospace applications, where guidance and control are implemented digitally on flight computers operating at fixed sampling intervals (ÅSTRÖM; WITTENMARK, 1997; STEVENS; LEWIS; JOHNSON, 2016).

2.7.1 Discrete-Time Plant Model

Consider the discrete-time linear time-invariant (LTI) system:

$$\begin{aligned}x_{k+1} &= Ax_k + Bu_k + w_k \\y_k &= Cx_k + v_k\end{aligned}\tag{2.24}$$

where $x_k \in \mathbb{R}^n$ is the state vector, $u_k \in \mathbb{R}^m$ the control input, and $y_k \in \mathbb{R}^p$ the measurement vector. The disturbances $w_k \sim \mathcal{N}(0, Q_w)$ and measurement noise $v_k \sim \mathcal{N}(0, R_v)$ are independent, zero-mean Gaussian sequences.

The performance index is a quadratic cost, defined by:

$$J = \lim_{N \rightarrow \infty} \mathbb{E} \left[\sum_{k=0}^N (x_k^\top Q x_k + u_k^\top R u_k) \right]\tag{2.25}$$

with $Q \succeq 0$ and $R \succ 0$ as design weights (ANDERSON; MOORE, 1979; ÅSTRÖM; WITTENMARK, 1997; STEVENS; LEWIS; JOHNSON, 2016).

2.7.2 LQR Component

Ignoring noise and assuming full-state availability, the optimal control minimizing J is given by the LQR solution in 2.5.2.

2.7.3 Kalman Filter Component

Since the true state x_k is usually not directly measurable, a state estimator is needed. The Kalman filter provides the minimum-variance linear estimate:

$$\hat{x}_{k+1|k} = (A_d - LC_d) \hat{x}_{k|k-1} + L y_k\tag{2.26}$$

with gain L expressed by (2.23).

2.7.4 LQG Controller by Separation Principle

By the separation principle, the regulator and estimator can be designed independently and then combined. The resulting LQG controller is:

$$u_k = -K \hat{x}_{k|k}\tag{2.27}$$

where $\hat{x}_{k|k}$ is provided by the Kalman filter and K is the LQR gain. This structure provides optimal regulation and estimation in the presence of Gaussian disturbances, and has become a standard design in aerospace guidance, navigation, and control (GNC) systems (STEVENS; LEWIS; JOHNSON, 2016).

The LQG controller inherits the optimality of LQR and Kalman filtering, yielding a powerful framework for digital control. However, Doyle famously pointed out that classical LQG lacks guaranteed robustness margins: “there are none” (DOYLE, 1978). This motivated the development of LTR and H_∞ methods, which explicitly shape open-loop frequency responses to enhance robustness. Despite this limitation, discrete-time LQG remains a cornerstone in modern aerospace applications, particularly for reusable launch vehicles, where precise estimation and optimal control under uncertainty are critical (ÅSTRÖM; WITTENMARK, 1997; STEVENS; LEWIS; JOHNSON, 2016).

2.8 Loop Transfer Recovery

The central principle of Loop Transfer Recovery (LTR) is to modify the estimator dynamics so that the open-loop transfer function of the LQG-controlled system recovers the gain and phase margins characteristic of the underlying LQR design. In practice, this is achieved by shaping the Kalman filter dynamics in such a way that the estimator approaches an identity mapping at low frequencies (STEVENS; LEWIS; JOHNSON, 2016; ZHOU; DOYLE; GLOVER, 1996; DOYLE; STEIN, 1979). A common heuristic consists in scaling the covariance matrix of the measurement noise R by a small parameter $\rho > 0$, while keeping the process noise covariance Q fixed. As $\rho \rightarrow 0$, the Kalman gain increases and the estimator bandwidth broadens, driving the estimator transfer function toward unity.

Formally, one considers the Kalman filter Riccati equation (2.22) and replaces R by ρR , with $\rho \ll 1$. In the limit $\rho \rightarrow 0$, the resulting estimator loop transfer $L_{est}(s)$ tends to one, so that the overall LQG open loop approaches the ideal LQR loop transfer:

$$L_{LQG}(s) \longrightarrow L_{LQR}(s).$$

This procedure ensures that the desirable robustness properties of LQR—namely satisfactory gain and phase margins—are approximately restored in the practical output-feedback implementation.

While conceptually elegant, this approach is not without limitations. Stevens (STEVENS; LEWIS; JOHNSON, 2016) emphasizes that LTR assumes the plant to be minimum-phase in order to recover the robustness margins effectively. For systems with non-minimum phase zeros, the attempt to push the estimator dynamics toward high bandwidth may lead to amplification of undesirable dynamics and to instability rather than margin

recovery. In such cases, the recourse to LTR may yield excessive control activity, degraded noise sensitivity, and ultimately poor closed-loop robustness. This observation highlights that, although LTR successfully embeds robustness into the LQG framework for many practical systems, its applicability is constrained when non-minimum phase behavior is present. For MIMO systems, the so-called nonminimumphase zeros correspond to the transmission zeros of the plant. In this context, the term nonminimumphase refers to those transmission zeros that lie in outside the unit disk (in the closedright half complex plane in continuous-time).

2.9 Robustness and Frequency-Domain Analysis

In control theory and aerospace applications, the concepts of *robustness* and *frequency-domain analysis* play a crucial role in ensuring system stability and performance under uncertainty. Robustness refers to the ability of a control system to maintain desired properties—such as stability, performance, and disturbance rejection—despite variations in system parameters, unmodeled dynamics, or external perturbations. This characteristic is particularly important in aerospace systems, where vehicles such as aircraft, launchers, and reusable rockets are exposed to highly variable operating conditions, significant uncertainties in aerodynamic coefficients, and severe external disturbances like wind gusts or turbulence.

The frequency-domain framework provides engineers with a powerful set of tools to evaluate and quantify robustness. Unlike time-domain methods, which focus on transient responses or steady-state errors, frequency-domain techniques use the spectral representation of signals and systems to capture stability margins, bandwidth, and the capacity of controllers to reject disturbances across different frequency ranges. Classical measures such as gain margin, phase margin, and sensitivity functions are defined in this domain, enabling a precise characterization of how uncertainties and perturbations affect system dynamics (STEVENS; LEWIS; JOHNSON, 2016; ÅSTRÖM; WITTENMARK, 1997).

2.9.1 Robustness in control

Robustness is often described as a measure of how well a system tolerates deviations from its nominal model. In practice, no mathematical model perfectly captures the real dynamics of an aircraft or launch vehicle. Structural flexibility, unmodeled high-frequency modes, nonlinear aerodynamic coefficients, and parameter variations due to fuel consumption or payload differences all contribute to uncertainty.

A robust system is one that remains stable and effective across these uncertainties. For instance, in a rocket descending with retro-propulsion, robustness ensures that changes in mass distribution or aerodynamic disturbances do not cause divergence in attitude or

trajectory. This property is not absolute but quantified: a system is considered robust if its stability margins and performance indices remain within acceptable limits for the range of expected variations.

Two perspectives on robustness are common:

- **Internal stability robustness:** the capacity of the closed-loop system to remain stable despite parameter variations.
- **Performance robustness:** the ability to satisfy performance criteria such as settling time, overshoot, or disturbance rejection under uncertainty.

These two dimensions must be simultaneously addressed, particularly in aerospace systems where safety-critical operations demand not only stability but also precise performance under harsh conditions.

2.9.2 Frequency-domain tools for robustness assessment

Frequency-domain analysis provides direct insights into robustness. The transfer functions of open-loop or closed-loop systems can be studied using Bode plots, Nyquist diagrams, or Nichols charts to reveal how disturbances at different frequencies affect system behavior.

Classical stability margins.

Two primary indicators of robustness are the gain margin and phase margin. The gain margin quantifies how much loop gain can increase before instability occurs, while the phase margin measures the additional phase lag required to push the system to the brink of instability. These margins are interpreted directly from frequency-response plots and offer engineers intuitive indicators of robustness. For aerospace systems, where uncertainties in aerodynamic coefficients can be significant, margins of at least 6 dB (gain) and 30–45° (phase) are often targeted (STEVENS; LEWIS; JOHNSON, 2016).

Sensitivity and complementary sensitivity functions.

The sensitivity function,

$$S_{\text{sen}}(e^{j\omega T_s}) = \frac{1}{1 + L(e^{j\omega T_s})} \quad (2.28)$$

where L is the loop transfer function, describes how the system responds to disturbances and model errors. A small sensitivity magnitude at low frequencies implies strong disturbance rejection, while high-frequency peaks indicate vulnerability to unmodeled dynamics.

The complementary sensitivity function,

$$T_{\text{com}}(e^{j\omega T_s}) = \frac{L(e^{j\omega T_s})}{1 + L(e^{j\omega T_s})} \quad (2.29)$$

instead reflects the ability to follow reference commands but also reveals susceptibility to measurement noise at high frequencies. Robust design requires a trade-off: suppressing sensitivity at low frequencies while avoiding excessive amplification at high frequencies. This balance is essential in launch vehicle guidance and control, where low-frequency robustness guarantees trajectory precision, and high-frequency attenuation prevents instability from structural vibrations.

Moreover, the sensitivity curve S_{sen} describes how the system reacts to disturbances injected at the output, whereas the complementary-sensitivity curve T_{com} shows how the system responds to variations in the reference signal (STEVENS; LEWIS; JOHNSON, 2016). From these curves, the maximum resonance peaks in the magnitudes of these functions over the indicated frequency ranges are evaluated as:

$$M_S \triangleq \max_{\omega} |S_{\text{sen}}(e^{j\omega T_s})| \quad (2.30)$$

$$M_T \triangleq \max_{\omega} |T_{\text{com}}(e^{j\omega T_s})| \quad (2.31)$$

Additionally, following Postlethwaite (1996), lower bounds for the gain margin (GM) and phase margin (PM) can be inferred from M_S and M_T as:

$$GM \geq \min \left\{ 20 \log_{10} \left(\frac{M_S}{M_S - 1} \right), 20 \log_{10} \left(\frac{M_T + 1}{M_T} \right) \right\} \quad (2.32)$$

$$PM \geq \left(\frac{180}{\pi} \right) \min \left\{ 2 \sin^{-1} \left(\frac{1}{2M_S} \right), 2 \sin^{-1} \left(\frac{1}{2M_T} \right) \right\} \quad (2.33)$$

These expressions provide a quantitative and consistent assessment of the robustness of the designed controllers (SILVEIRA et al., 2025; SEBORG; EDGAR; MELLICHAMP, 2003).

Frequency-domain analysis offers a rich and intuitive framework for this purpose. Classical indicators such as gain and phase margins, sensitivity functions, and more advanced tools like μ -analysis provide quantitative measures of how close a system is to instability and how well it tolerates uncertainty. These tools not only diagnose weaknesses but also directly inform the design of robust controllers, from loop-shaping of PID laws to advanced H_{∞} synthesis.

In the broader context of aerospace, robustness and frequency-domain analysis form the backbone of modern guidance and control. They bridge theory and practice, ensuring that vehicles ranging from commercial aircraft to reusable launchers can perform reliably in the face of uncertainties. The synergy between robust control theory and frequency-domain tools continues to shape the future of aerospace systems, enabling safer, more efficient, and more autonomous flight.

2.10 Scheduling

In trajectory-dependent flight regimes, a single fixed controller is often too conservative. *Gain scheduling* addresses this by designing local LTI controllers (or gains) at representative operating points and interpolating them online as a function of a scheduling vector ρ (e.g., altitude, Mach, dynamic pressure). When the rate of change $\dot{\rho}$ is small relative to the closed-loop bandwidth, the system behaves close to the “frozen” LTI models and scheduled gains can deliver near-local optimal performance across a wide envelope. The practical challenge then shifts to **how** gains are interpolated between the scheduled points, because the choice of interpolation directly affects smoothness (bumplessness), transient excitation, and robustness margins.

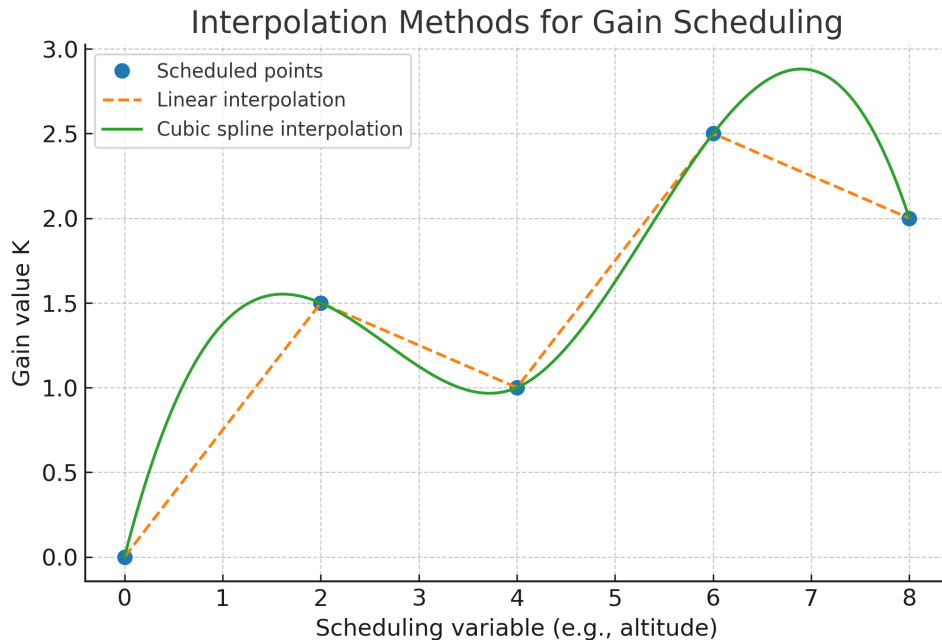


Figure 2.6 – Illustration of scheduled points and two interpolation laws for a scalar gain $K(\rho)$: piecewise linear (dashed) and cubic spline (solid).

Source: Open AI

Figure 2.6 illustrates typical behaviours for a scalar gain $K(\rho)$: piecewise linear interpolation is simple and predictable but introduces slope discontinuities at knots; cubic

splines create C^2 continuity and often smaller steady interpolation error, but may overshoot and should be constrained (e.g., monotone Hermite/PCHIP) when K must remain positive or monotone. In two or more dimensions, bilinear/tri-linear (or barycentric on simplices) blends are common because they are fast and preserve the local convex hull of the corner values.

Table 2.1 summarizes compact formulations you can reuse in implementation. In practice, three guidelines help preserve robustness when scheduling: (i) ensure continuity of $K(\rho)$ and, when possible, of its first derivative to avoid excitation of lightly damped modes; (ii) filter or rate-limit measured ρ and enforce clamping outside the grid; (iii) place grid points where the plant varies most (changes in M , q), and check margins along fast sweeps of ρ to bound “interpolation uncertainty”.

Table 2.1 – Common interpolation methods for gain scheduling

Method	Formulation
Linear (1D)	$K(x) = K_i + \frac{x - x_i}{x_{i+1} - x_i} (K_{i+1} - K_i)$ <p>where x is the scheduling variable (altitude, Mach, etc.), and K_i, K_{i+1} are the gains at the neighboring operating points.</p>
Bilinear (2D)	$K(x, y) = \sum_{m=0}^1 \sum_{n=0}^1 w_{mn} K_{i+m, j+n}$ <p>with w_{mn} the weights computed from the relative distances in both directions (x and y), and $K_{i,j}$ the gains at the grid points.</p> <p>Smooth interpolation defined by</p>
Cubic spline (1D)	$K(x) = a_0 + a_1(x - x_i) + a_2(x - x_i)^2 + a_3(x - x_i)^3,$ <p>with coefficients a_k determined from the points and boundary derivatives, ensuring continuity of the first and second derivatives.</p>

Remark.

If positivity/monotonicity of K is critical, prefer shape-preserving splines (PCHIP/monotone Hermite) over cubic polynomials, and combine with bumpless transfer and rate limits on ρ for smooth closed-loop behaviour.

2.11 Summary

This chapter has established the theoretical groundwork for robust control of RLVs during aerodynamic descent. A CALLISTO-inspired 6-DoF model with Mach-, angle-of-

attack-, and control-deflection-dependent aerodynamics under realistic disturbances was formulated, leading to a state-space representation with local linearizations suitable for controller design.

On top of this modeling layer, the LQG architecture was introduced, combining LQR optimal regulation with Kalman filtering for state estimation, together with Loop Transfer Recovery (LTR) to regain LQR-like loop shapes and classical robustness margins. Frequency-domain tools (sensitivity and complementary sensitivity and stability margins) and scheduling with interpolation across synthesis points were outlined as the main instruments to handle large envelope variations from supersonic to subsonic regimes.

3 Methodology

3.1 Initial Considerations

The methodology adopted in this work is designed to address the specific challenges associated with the aerodynamic descent of RLVs. This phase of flight is characterized by strong nonlinearities, rapid parameter variations, and significant aerodynamic uncertainties, all of which must be accommodated while ensuring precise trajectory tracking and robust stability margins.

The proposed approach is structured to progressively transform the high-fidelity nonlinear vehicle model into a form suitable for advanced control design, specifically the Linear Quadratic Gaussian with Loop Transfer Recovery LQG/LTR framework. The process begins with the definition of the vehicle dynamics, atmospheric conditions, and mission profile, followed by the selection of representative operating points for linearization. These linearized models serve as the foundation for controller synthesis and gain scheduling across the flight envelope.

The methodology is organized into six main stages:

1. **Synthesis Point Identification for Scheduling:** Procedure for extracting and validating linearization points along the nominal descent trajectory, ensuring representative coverage of aerodynamic regimes.
2. **Linearization and Discretization:** Extraction of linear model by synthesis points and discretization for digital controllers design.
3. **Model Augmentation and Pre-compensator Design:** Inclusion of integral action for offset elimination, actuator dynamics modeling, and pre-compensator shaping to achieve desired frequency-domain characteristics.
4. **Proposed Controller Architecture:** Definition of the control structure, including state feedback, optimal estimation, Loop Transfer Recovery augmentation, and the integration of these elements into a unified framework suitable for gain scheduling.
5. **Summary and Conclusions:** Consolidation of methodological choices and preparation for the simulation and validation phases.

This structure ensures a logical flow from problem definition to robust control law synthesis, bridging the gap between high-fidelity aerodynamic modeling and implementable guidance and control solutions for RLV atmospheric descent.

3.2 Synthesis Point Identification for Scheduling

In Sagliano et al. (2023), the selection of linearization points for gain-scheduling is performed by sampling the reference trajectory at 15 operating conditions defined by a single, monotonically varying scheduling parameter: the altitude. This choice avoids ambiguity between different physical states that might share other parameters (e.g., Mach number), ensures each point represents a unique flight condition, and allows accurate onboard measurement without model-based estimation. At each chosen altitude, the nonlinear six-degree-of-freedom equations of motion are analytically linearized to obtain a set of LTI models for controller synthesis and subsequent interpolation along the descent trajectory.

Instead of heuristically selecting the operating points, we propose an analysis of the reference trajectory based on detecting the inflection points of the curve. This method, which involves computing the second derivative point by point, identifies the inflection point(s) under the assumption that these coincide with critical maneuver points.

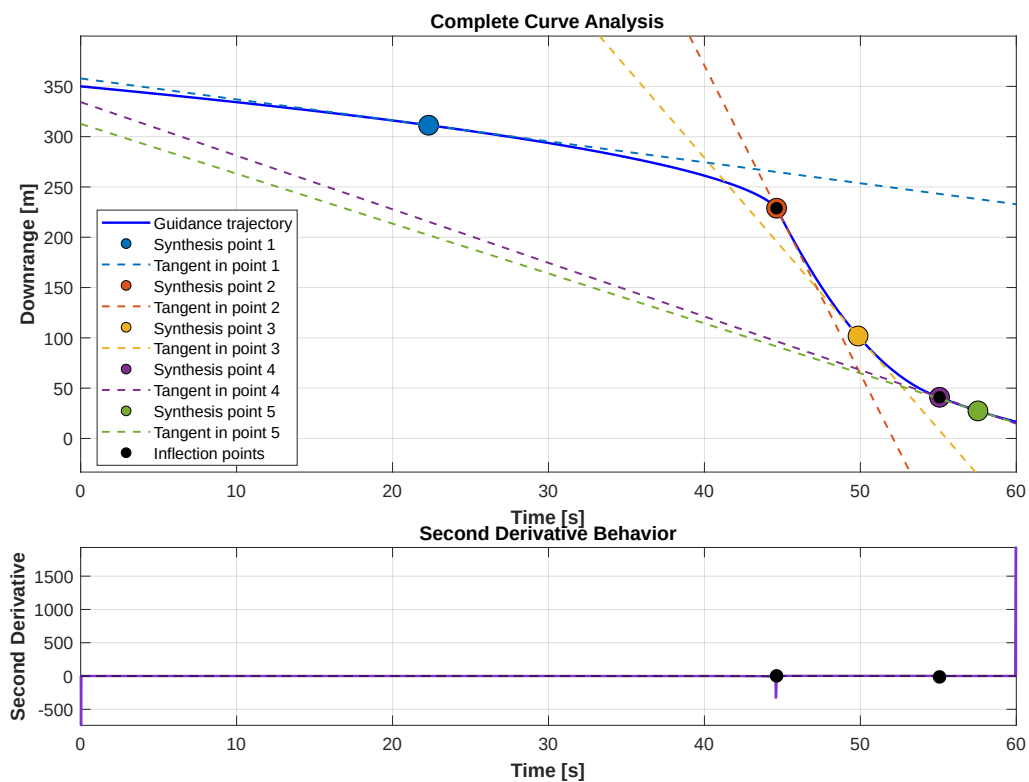


Figure 3.1 – Proposed operation points around a suggested trajectory for back to lunch site scenario by right.

Source: Author

To mitigate the nonlinear effects of the system, we add supplementary points

following a rule we call the ***Tangent Rule***. The procedure consists of a geometric-visual analysis of the trajectory. In the first stage, supplementary points are introduced between the endpoints and the predefined reference points (notably inflection points), with the objective of distributing them at approximately equal intervals along the curve. In cases where multiple inflection points are present, the same principle is applied recursively to ensure a uniform distribution of intermediate points. In the second stage, tangents are constructed to the trajectory at each of the selected points. For every pair of consecutive points, the corresponding tangents are extended, and their intersection is determined. These intersection points are subsequently projected onto the original curve, thereby generating additional operating points whenever the spacing between the existing ones is considered too large. This construction provides a refined set of reference points that enhances the resolution and accuracy of the trajectory analysis.

The Figure 3.1 reveals the operating points obtained from the analysis of the reference trajectory employed in this dissertation through the proposed method. Based on this analysis, five representative operating points were selected as the basis for the controller design.

3.3 Model Augmentation and Pre-compensator Design

To ensure a unitary gain in steady state, a static gain type pre-compensator is a better option. This adjusts the reference with the inverse of the DC gain.

$$\mathbf{P}_c^* = \left(\mathbf{C}_d (\mathbf{I}_{n_x \times n_x} - \mathbf{A}_d)^{-1} \mathbf{B}_d + \mathbf{D}_d \right)^{-1} \quad (3.1)$$

The matrix \mathbf{B}_d is transformed in the pre-compensated as follows:

$$\mathbf{B} = \mathbf{P}_c \times \mathbf{B}_d \quad (3.2)$$

Including discrete integral action (incremental action) to model-based digital control systems, requires necessary augmented design model. An alternative is to work with the discrete control variation. That is,

$$\Delta \mathbf{u}(k) = \mathbf{u}(k) - \mathbf{u}(k - 1) \quad (3.3)$$

such that the control signal that is in fact applied at the process entrance is

* *System dimensions:* n_x number of states,
 n_y number of outputs,
 n_u number of inputs.

$$\mathbf{u}(k) = \mathbf{u}(k-1) + \Delta\mathbf{u}(k) \quad (3.4)$$

Then, to obtain the augmented model by the discrete difference operator, $\Delta = 1 - z^{-1}$, state and output equations are modified to

$$\begin{aligned} \Delta\mathbf{x}(k+1) &= \mathbf{A}_d\Delta\mathbf{x}(k) + \mathbf{B}\Delta\mathbf{u}(k) \\ \Delta\mathbf{y}(k+1) &= \mathbf{C}_d\Delta\mathbf{x}(k+1) \end{aligned} \quad (3.5)$$

and a new augmented state vector is defined as:

$$\mathbf{x}_a(k) = \begin{bmatrix} \mathbf{y}_a(k) \\ \Delta\mathbf{x}(k) \end{bmatrix}, \quad (3.6)$$

Since $\Delta\mathbf{y}(k+1) = \mathbf{y}(k+1) - \mathbf{y}(k)$, then the output and discrete state change equations can be written as:

$$\begin{aligned} \mathbf{y}(k+1) &= \mathbf{y}(k) + \mathbf{C}_d\mathbf{A}_d\Delta\mathbf{x}(k) + \mathbf{C}_d\mathbf{B}\Delta\mathbf{u}(k) \\ \Delta\mathbf{x}(k+1) &= \mathbf{A}_d\Delta\mathbf{x}(k) + \mathbf{B}_d\Delta\mathbf{u}(k) \end{aligned} \quad (3.7)$$

The augmented model obtained is the following:

$$\begin{aligned} \mathbf{x}_a(k) &= \mathbf{A}_a\mathbf{x}_a(k-1) + \mathbf{B}_a\Delta\mathbf{u}(k-1) \\ \mathbf{y}_a(k) &= \mathbf{C}_a\mathbf{x}_a(k) \end{aligned} \quad (3.8)$$

with:

$$\mathbf{A}_a = \begin{bmatrix} \mathbf{I}_{n_y \times n_y} & (\mathbf{C}_d\mathbf{A}_d)_{n_y \times n_x} \\ \mathbf{0}_{n_x \times n_y} & (\mathbf{A}_d)_{n_x \times n_x} \end{bmatrix}; \mathbf{B}_a = \begin{bmatrix} (\mathbf{C}_d\mathbf{B})_{n_y \times n_u} \\ \mathbf{B}_{n_x \times n_u} \end{bmatrix}; \mathbf{C}_a = \begin{bmatrix} \mathbf{I}_{n_y \times n_y} & \mathbf{0}_{n_y \times n_x} \end{bmatrix} \quad (3.9)$$

For trajectory tracking during the rocket's aerodynamic descent phase in a noisy environment, the incremental servo controller designed, based on LQG/LTR methodology, aims at providing robust control, applying the incremental control law described in 3.4 with:

$$\begin{aligned} \Delta\mathbf{u}(k) &= \mathbf{K}_{servo}\mathbf{r}(k) - \mathbf{K}_{lqr}\bar{\mathbf{x}}_a(k) \\ \mathbf{K}_{servo} &= \mathbf{K}_{lqr}(1 : n_y) \end{aligned} \quad (3.10)$$

Standard solutions for both the LQR problem and its dual counterpart, the Kalman state estimation problem, are readily available in modern control system design environ-

ments, ranging from proprietary to open-source platforms. Consequently, the implementation challenge primarily reduces to proper tuning of the cost function minimization parameters

$$J_{LQR} = \sum_{k=0}^{\infty} \mathbf{x}_a^T(k) \mathbf{Q} \mathbf{x}_a(k) + \Delta \mathbf{u}^T(k) \mathbf{R} \Delta \mathbf{u}(k) \quad (3.11)$$

where \mathbf{Q} and \mathbf{R} are used to weight the optimization with regards to the state vector and the control increment.

3.4 Proposed Controller Architecture

In line with recent GNC developments for reusable first stages, it was adopted a stratified Guidance-Navigation-Control architecture tailored to aerodynamic descent and precise landing. The **Guidance** module computes flyable reference profiles (attitude commands, body-velocity references, and thrust direction/magnitude where applicable) using a successive-convexification strategy solved in real time, which naturally enforces state/control constraints and accommodates mission objectives such as time and fuel usage (OLIVEIRA; LAVAGNA, 2024). The **Navigation** module delivers time-consistent state estimates (attitude, angular rates, lateral body-velocity components, etc.) via Kalman filtering, providing to the controller both the feedback signals and the scheduling variables required to span multiple flight regimes.

The **Control** module is based on a LQG regulator *rendered robust* through LTR. Concretely, an LQR loop-shape is first specified at the plant output; then, by appropriately inflating the process-noise weighting in the Kalman filter, the recovered open-loop approaches the target LQR loop, thereby improving classical stability margins and disturbance rejection while preserving optimal tracking and noise-attenuation properties. This LQG/LTR design is carried out at a set of frozen operating points along the descent and then gain-scheduled (e.g., versus Mach, dynamic pressure, or altitude) to cover the large variations in controllability and actuator authority typical of aerodynamic return trajectories (SAGLIANO et al., 2023; OLIVEIRA; LAVAGNA, 2024).

Overall, this unified GNC stack, *convex guidance in the loop, Kalman navigation, and multi-regime LQG/LTR control*, aims to combine high tracking accuracy under constraints with explicit robustness margins across flight regimes and actuator configurations, as advocated in recent CALLISTO-related and RLV descent studies (SAGLIANO et al., 2023; OLIVEIRA; LAVAGNA, 2024).

The Figure 3.2 summarizes the CALLISTO-inspired GNC structure adopted here: the Guidance module provides feasible reference profiles $\mathbf{y}_{\text{ref}} = [\mathbf{r}^\top, \mathbf{v}^\top, \mathbf{e}\mathbf{a}^\top, \boldsymbol{\omega}^\top]^\top$ ($\mathbf{e}\mathbf{a} = (\phi, \theta, \psi)$ Euler angles in the body-fixed). A Kalman-based Navigation stage (embedded in

the controller block) delivers consistent state estimates to a LQG regulator *rendered robust* via LTR, which shapes the recovered open-loop toward an LQR, like loop with comfortable classical margins. The Plant collects the 6-DoF rigid-body and aerodynamic dynamics, actuator dynamics (e.g., fin/TVC), and exogenous wind disturbances \mathbf{d}_w . Control inputs are $\mathbf{u} = [\delta_\phi, \delta_\theta, \delta_\psi]^\top$. This layout matches recent RLV descent studies advocating convex guidance in the loop and robust LQG/LTR tracking across flight regimes.

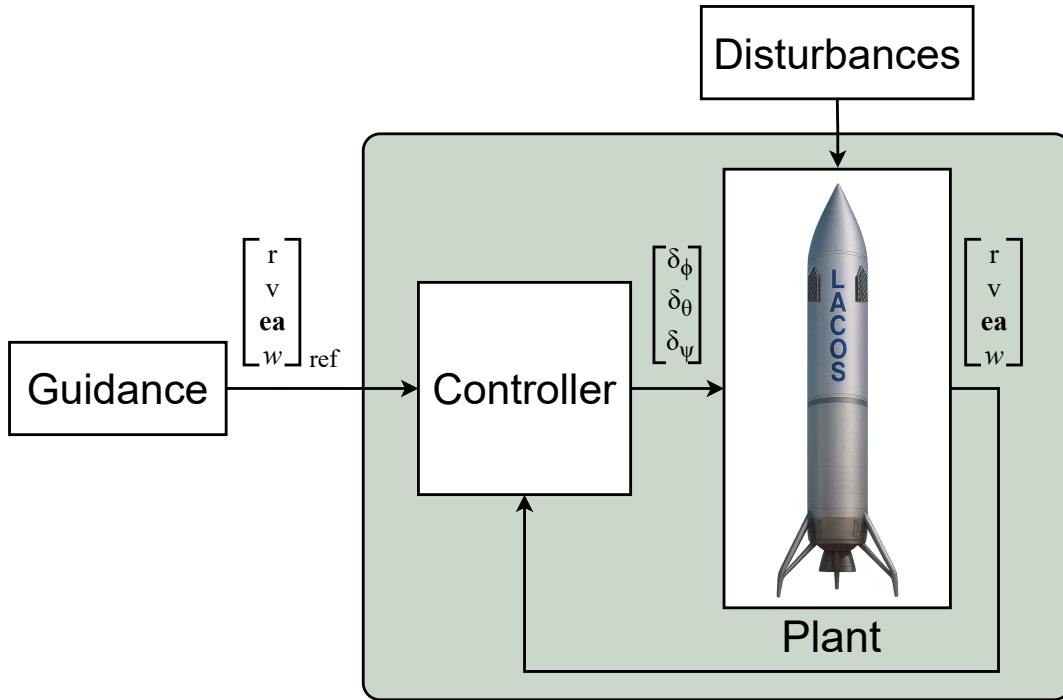


Figure 3.2 – CALLISTO-based RLV control loop for the aerodynamic descent subject to wind disturbance.

Source: Author

Structure and rationale: The Figure 3.3 shows the discrete-time implementation of the LQG controller robustified via LTR, used identically on roll and on the pitch/yaw pair (the vehicle’s quasi-symmetric geometry enables a common design and joint analysis for pitch and yaw). Guidance provides the reference \mathbf{y}_{ref} , the controller forms the tracking error, and applies a pre-compensator K_s together with the state-feedback law

$$u(k) = -\mathbf{K}_{lqr} \hat{\mathbf{x}}(k) + \mathbf{K}_{servo} \mathbf{r}(k) \quad (3.12)$$

on the *augmented* model (A_a, B_a, C_a) , which includes an integrator for offset-free tracking. A Kalman observer with gain L reconstructs $\hat{\mathbf{x}}(k)$ from the measured output, while the unit delays z^{-1} emphasize the sampled implementation. The **System** block collects the linearized plant matrices $[A, B, C]$ and the exogenous wind disturbance; the plant output \mathbf{y} is fed back to both controller and observer. In the **LTR** step, the filter weightings are

tuned (progressively inflating the process-noise level relative to the measurement-noise level) so that the recovered open-loop approaches the target LQR loop-shape at the plant output, thereby improving classical gain/phase margins and disturbance rejection without sacrificing the tracking performance imposed by guidance. The CALLISTO-inspired model used here is kept *decoupled* on each axis to streamline frequency-domain analysis and per-axis validation; a subsequent gain-scheduling (e.g., with Mach, dynamic pressure, or altitude) maintains performance over the full descent envelope.

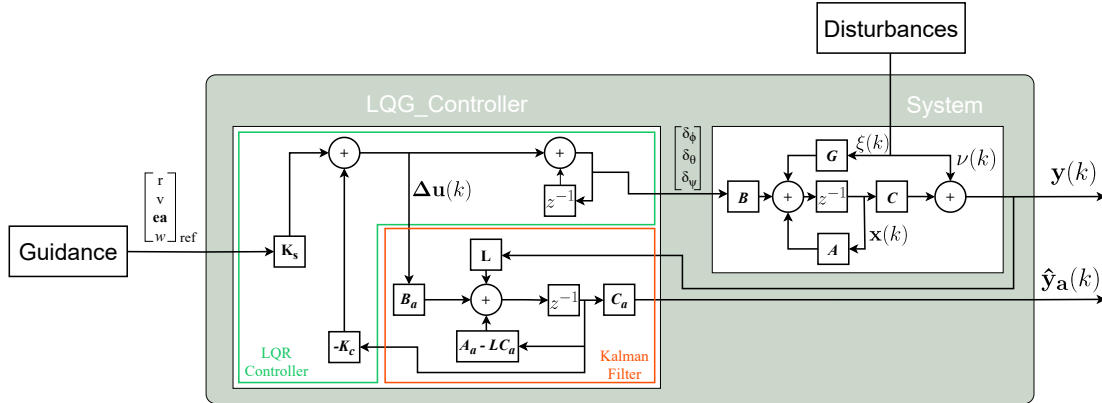


Figure 3.3 – Block diagram of the LQG control loop for the RLV aerodynamic descent.

Source: Author

3.4.1 Atmosphere and Disturbances

Thermodynamic properties follow the International Standard Atmosphere (ISA): density $\rho(h)$, temperature $T(h)$, and speed of sound $a(h)$. Winds combine (i) a deterministic mean profile and (ii) stochastic turbulence (Dryden or von Kármán) generated by shaping filters (STEVENSON; LEWIS; JOHNSON, 2016). Disturbances act on translation through the air-relative velocity \mathbf{v}_{air} and on rotation through dynamic-pressure-scaled derivatives, ensuring consistency in robustness studies (SAGLIANO et al., 2023; OLIVEIRA; LAVAGNA, 2024).

We formulate the wind field in three steps to keep a clear separation between (a) the background (mean) profile, (b) coordinate consistency with the vehicle attitude, and (c) stochastic turbulence consistent with standard spectra.

Step 1. Mean horizontal wind in local North–East–Down frame (NED), with smooth piece-wise law.

Let $\mathcal{F}_n = \{\mathbf{e}_N, \mathbf{e}_E, \mathbf{e}_D\}$ NED and z the above ground level (AGL) altitude. The mean speed $U(z)$ is synthesized by smoothly blending:

(i) Log law (neutral):

$$\begin{aligned} U_{\log}(z) &= \frac{u_*}{\kappa} \ln\left(\frac{z+z_0}{z_0}\right) \\ u_* &= \frac{\kappa U_{100}}{\ln\left(\frac{100+z_0}{z_0}\right)} \end{aligned} \quad (3.13)$$

(ii) Power law (anchored at 100 m):

$$\begin{aligned} U_{\text{pow}}(z) &= U_{100} \left(\frac{\max(z,100)}{100}\right)^\alpha \\ U_{100} &= \frac{u_*}{\kappa} \ln\left(\frac{100+z_0}{z_0}\right) \end{aligned} \quad (3.14)$$

(iii) LLJ / mid-jet transition (1–3 km):

$$U_{\text{tran}}(z) = U_{\text{LLJ}} + w_3(z) (U_{\text{mid}} - U_{\text{LLJ}}) \quad (3.15)$$

(iv) Subtropical jet (Gaussian core):

$$U_{\text{jet}}(z) = U_{\text{jet,pk}} \exp\left(-\left(\frac{z-z_{\text{jet}}}{\sigma}\right)^2\right) \quad (3.16)$$

Smooth C^1 weights $w_{12}, w_{23}, w_{34} \in [0, 1]$ blend the stages across altitude bands $(z_{12}^1, z_{12}^2), (z_{23}^1, z_{23}^2), (z_{34}^1, z_{34}^2)$:

$$\begin{aligned} U_{12} &= (1 - w_{12})U_{\log} + w_{12}U_{\text{pow}} \\ U_{23} &= (1 - w_{23})U_{12} + w_{23}U_{\text{tran}} \\ U(z) &= (1 - w_{34})U_{23} + w_{34}U_{\text{jet}} \end{aligned} \quad (3.17)$$

Seasonal parameters ($\alpha, z_{\text{jet}}, U_{\text{jet,pk}}, U_{\text{LLJ}}$) change with the month. The mean wind direction uses the *TO* convention:

$$\psi_w^{\text{TO}}(z) = \underbrace{\psi_{\text{base}} + 180 \mathbb{I}_{\text{FROM}}}_{\text{handles dir_is_from}} + \underbrace{\Delta\psi_{\text{max}} \left[1 - \exp\left(-\left(\frac{z}{z_v}\right)^p\right)\right]}_{\text{Ekman veering}} + \underbrace{\varepsilon \sin\left(\frac{z}{\ell}\right)}_{\text{small shear osc.}^1} \quad (3.18)$$

so the mean NED components are

$$\bar{\mathbf{w}}_n(z) = \left[U(z) \cos \psi_w^{\text{TO}}, U(z) \sin \psi_w^{\text{TO}}, 0 \right]^\top. \quad (3.19)$$

Step 2. Frame-consistent projection and air-relative velocity. With attitude $\boldsymbol{\eta}$, the DCM $C_{b \leftarrow n}(\boldsymbol{\eta})$ gives $\bar{\mathbf{w}}_b = C_{b \leftarrow n} \bar{\mathbf{w}}_n$. The air-relative velocity used in aerodynamics is then $\mathbf{v}_{\text{air}} = \mathbf{v}_b - (\bar{\mathbf{w}}_b + \tilde{\mathbf{w}}_b)$, where $\tilde{\mathbf{w}}_b$ is the zero-mean turbulent component (next step).

¹ oscillation

Step 3. Stochastic turbulence by shaping filters (optional in this profile). Gusts $\tilde{\mathbf{w}}_n = [u_g, v_g, w_g]^T$ in NED are generated via Dryden or von Kármán filters, with altitude-scheduled intensities $\sigma_{u,v,w}(z)$ and scale lengths $L_{u,v,w}(z)$: $\dot{\mathbf{x}}_g = A_g(z)\mathbf{x}_g + B_g(z)\boldsymbol{\xi}$, $\tilde{\mathbf{w}}_n = C_g(z)\mathbf{x}_g + D_g(z)\boldsymbol{\xi}$ (discretized at T_s). They are then rotated to body axes, $\tilde{\mathbf{w}}_b = C_{b \leftarrow n}\tilde{\mathbf{w}}_n$, and added to the mean field above.

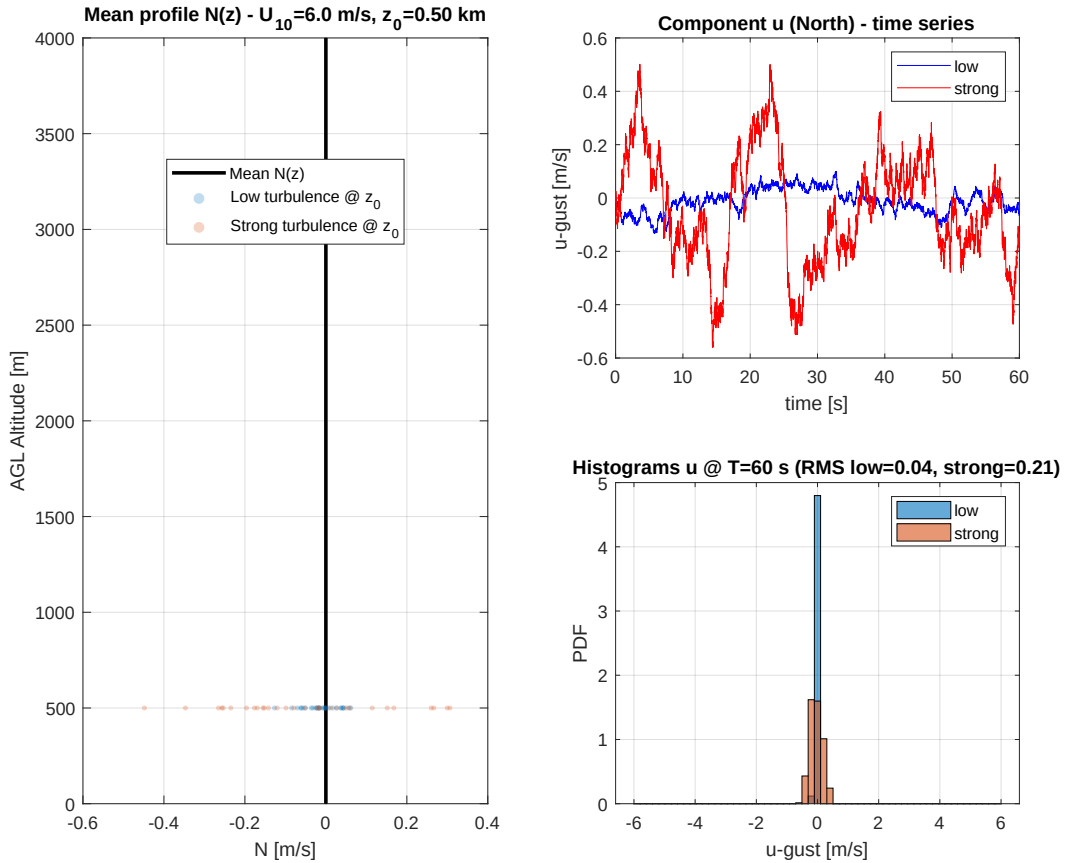


Figure 3.4 – North-wind profile and turbulence characterization. (a) Mean northward wind $N(z)$. (b) Time series of the u (north) component. (c) Histograms/PDFs of u -gust (light and strong).

Source: Author

In Figure 3.4, the left panel gives the mean vertical wind profile $N(z)$ versus AGL altitude for $U_{10} = 6$ m/s and $z_0 = 0.50$ km; the middle panel contrasts time series of the u (north) component over 060s for light versus strong turbulence; the right panel compares probability density function (PDF)s (histograms) of u -gust at $T = 60$ s, where the strong case exhibits a wider, flatter distribution (RMS 0.21 m/s) than the light case (RMS 0.04 m/s).

3.5 Linearization and Discretization

3.5.1 Numerical Considerations

For digital control, each linear model along the descent is discretized with a zero-order hold (ZOH). In practice we use MATLAB's `c2d` on the state-space realization $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$, $\mathbf{y} = \mathbf{C}\mathbf{x} + \mathbf{D}\mathbf{u}$ at the sampling period $T_s = 0.001$ s (1 kHz):

$$\{A_d, B_d, C_d, D_d\} = \text{c2d}(\{A, B, C, D\}, T_s, \text{'zoh'}). \quad (3.20)$$

The choice $T_s = 1$ ms satisfies the design rule $T_s \leq 1/(10f_{\max})$ with margin, where f_{\max} is the highest scheduled closed-loop target bandwidth. ZOH preserves the continuous loop shape closely in the frequency range of interest and avoids the additional warping induced by bilinear/Tustin transforms when strong loop-shaping (e.g., LTR) is used (MATHWORKS, 2024; OPPENHEIM; SCHAFER, 2010).

Time-varying schedule.

Discretization is performed *per operating point* (Mach/altitude/dynamic-pressure grid), yielding a bank $\{A_d(k), B_d(k), C_d(k), D_d(k)\}$ that is interpolated (or switched with hysteresis) during simulation and synthesis. To maintain numerical consistency, the same $T_s = 1$ ms is used for plant, controller, and estimators.

Observer/Kalman implementation.

When using LQG/LTR, the estimator runs at the same T_s and uses the discrete plant (A_d, B_d, C_d, D_d) together with process/measurement covariances (Q_d, R_d) associated. The Kalman gain L is obtained from the standard discrete-time recursions show in 2.7.3 and, when needed, is gain-scheduled jointly with the controller gains (ÅSTRÖM; WITTENMARK, 1997; ANDERSON; MOORE, 1979).

Quaternion and aerodynamic numerics.

Quaternion kinematics are integrated at $T_s = 1$ ms with unit-norm renormalization every step to prevent drift. Aerodynamic interpolation uses the *air-relative* variables (α, β, M, Re) computed from \mathbf{v}_{air} (wind-corrected), not the ground-relative velocity \mathbf{v} , ensuring the consistency of look-ups across wind conditions (STEVENS; LEWIS; JOHNSON, 2016; SAGLIANO et al., 2023; OLIVEIRA; LAVAGNA, 2024).

Minimal code stub (for each grid point).

```

Ts = 1e-3; % 1 kHz
sysc = ss(A,B,C,D); % continuous frozen model
sysd = c2d(sysc, Ts, 'zoh'); % for scheduling bank
Ad = sysd.A; Bd = sysd.B; Cd = sysd.C; Dd = sysd.D;

```

3.6 Summary

This chapter established control-oriented workflow for the aerodynamic descent of reusable launch vehicles. Starting from a high-fidelity nonlinear description, we defined a disciplined path to LQG/LTR synthesis and scheduling, with explicit numerical and environmental assumptions:

- **Problem framing and operating points.** We grounded the methodology in the multi-regime, highly perturbed nature of aerodynamic descent and selected representative synthesis points along the reference trajectory. Beyond altitude-based sampling (SAGLIANO et al., 2023), we introduced a trajectory-geometry method (inflection and Tangent Rule) to refine point placement when large curvature or maneuver intensity suggests denser coverage.
- **Linearization and discretization.** Each model was linearized about its operating condition and discretized via ZOH using MATLAB’s `c2d` at a unified sampling period $T_s = 1$ ms. This choice respects the guideline $T_s \leq 1/(10f_{\max})$ and keeps plant, controller, and observer numerics coherent across the gain-scheduled bank.
- **Augmentation and pre-compensator.** We adopted an incremental (difference-operator) servo structure, added integral action for offset-free tracking, and applied a static pre-compensator to normalize DC gain. The resulting augmented model provides a clean interface for LQR/LQG design in discrete time.
- **Controller architecture.** The control core is an LQG regulator *rendered robust* via Loop-Transfer Recovery. By shaping the recovered open-loop toward an LQR target, LTR improves classical margins without sacrificing noise attenuation. An approach aligned with recent GNC practice for RLV descent (SAGLIANO et al., 2023; OLIVEIRA; LAVAGNA, 2024).
- **Atmosphere and disturbances.** The environment model follows ISA thermodynamics and a three-step wind formulation: smoothly blended mean speed/direction in NED, consistent projection to body axes, and optional Dryden/von Kármán turbulence via shaping filters. Disturbance coupling is enforced through air-relative variables, ensuring physically consistent robustness studies (STEVENS; LEWIS; JOHNSON, 2016).
- **Implementation notes.** Quaternion renormalization, anti-alias prefiltering, and discrete anti-windup were specified to preserve stability margins and tracking across

regimes; scheduling (Mach, dynamic pressure, altitude) uses interpolation or guarded switching with hysteresis.

Overall, the chapter bridges high-fidelity vehicle/atmosphere modeling and an implementable, gain-scheduled LQG/LTR design. It delivers the artifacts (discrete augmented models, pre-compensator, estimator/controller gains, and scheduling logic) required for the simulation and validation campaigns presented next (SAGLIANO et al., 2023; OLIVEIRA; LAVAGNA, 2024; STEVENS; LEWIS; JOHNSON, 2016).

4 Results

This chapter reports the quantitative evaluation of the proposed control strategies for the aerodynamic descent of RLVs. We first detail the simulation setup, test scenarios and performance indices, then compare baseline and proposed controllers, assess robustness in the frequency domain, and conclude with stress testing, scheduling aspects and a synthesis of findings.

4.1 Simulation Setup and Assumptions

4.1.1 Models and Implementation

The plant considered here is open-loop unstable, with dynamics that diverge unless feedback is applied. In such conditions, classical frequency-domain tools (Bode, Nyquist, Nichols) must be interpreted with care: raw open-loop margins are not meaningful until a stabilizing controller is in place and the internal closed-loop signals are bounded. Accordingly, we first synthesize a stabilizing feedback law to relocate the dominant dynamics to a stable regime; only then do we proceed with the frequency-domain assessment, examining loop shape, bandwidth, gain/phase margins, and the sensitivity/complementary-sensitivity trade-offs, to quantify performance and robustness.

The plant was linearized at five synthesis points, obtaining five distinct LTI models. For concision, only two representative discrete-time state-space realizations was presented (the other three are analogous and omitted).

$$\Sigma_2 : \quad x_{k+1} = A_2 x_k + B_2 u_k, \quad y_k = C_2 x_k + D_2 u_k,$$

$$\Sigma_4 : \quad x_{k+1} = A_4 x_k + B_4 u_k, \quad y_k = C_4 x_k + D_4 u_k.$$

roll model 2

$$A_{R,2} = \begin{bmatrix} 1 & 0.00099956 \\ 0 & 0.99912 \end{bmatrix} \quad B_{R,2} = \begin{bmatrix} 0.000178168 \\ 0.356284 \end{bmatrix} \quad C_{R,2} = \begin{bmatrix} 1 & 0 \end{bmatrix} \quad D_{R,2} = \begin{bmatrix} 0 \end{bmatrix}$$

roll model 4

$$A_{R,4} = \begin{bmatrix} 1 & 0.000999295 \\ 0 & 0.998591 \end{bmatrix} \quad B_{R,4} = \begin{bmatrix} 0.00018233 \\ 0.364659 \end{bmatrix} \quad C_{R,4} = \begin{bmatrix} 1 & 0 \end{bmatrix} \quad D_{R,4} = \begin{bmatrix} 0 \end{bmatrix}$$

Downrange model 2

$$A_{P,2} = \begin{bmatrix} 1 & 0.00099956 & 0 & 0 \\ 0 & 0.99912 & 0 & 0 \\ 0 & 0 & 1 & 0.00099956 \\ 0 & 0 & 0 & 0.99912 \end{bmatrix} \quad B_{P,2} = \begin{bmatrix} 3.17978e-05 \\ 0.0633581 \\ 0.000126718 \\ 0.253432 \end{bmatrix}$$

$$C_{P,2} = [1 \ 0 \ 0 \ 0] \quad D_{P,2} = [0]$$

Downrange model 4

$$A_{P,4} = \begin{bmatrix} 1 & 0.000999295 & 0 & 0 \\ 0 & 0.998591 & 0 & 0 \\ 0 & 0 & 1 & 0.000999295 \\ 0 & 0 & 0 & 0.998591 \end{bmatrix} \quad B_{P,4} = \begin{bmatrix} 3.54986e-05 \\ 0.0709923 \\ 0.000141994 \\ 0.283969 \end{bmatrix}$$

$$C_{P,4} = [1 \ 0 \ 0 \ 0] \quad D_{P,4} = [0]$$

Crossrange model 2

$$A_{Y,2} = \begin{bmatrix} 1 & 0.00099956 & 0 & 0 \\ 0 & 0.99912 & 0 & 0 \\ 0 & 0 & 1 & 0.00099956 \\ 0 & 0 & 0 & 0.99912 \end{bmatrix} \quad B_{Y,2} = \begin{bmatrix} 2.50939e-05 \\ 0.0501876 \\ 0.000100362 \\ 0.20075 \end{bmatrix}$$

$$C_{Y,2} = [1 \ 0 \ 0 \ 0] \quad D_{Y,2} = [0]$$

Crossrange model 4

$$A_{Y,4} = \begin{bmatrix} 1 & 0.000999295 & 0 & 0 \\ 0 & 0.998591 & 0 & 0 \\ 0 & 0 & 1 & 0.000999295 \\ 0 & 0 & 0 & 0.998591 \end{bmatrix} \quad B_{Y,4} = \begin{bmatrix} 2.78718e-05 \\ 0.0557324 \\ 0.000111487 \\ 0.22293 \end{bmatrix}$$

$$C_{Y,4} = [1 \ 0 \ 0 \ 0] \quad D_{Y,4} = [0]$$

After extracting the axis-wise state-space matrices at the selected synthesis points, the 3-DoF full model (Roll–Pitch–Yaw) is rebuilt by stacking the three SISO (or per-axis) subsystems on the diagonal. In MATLAB, this is done with `blkdiag`, which preserves each axis independently while forming one global LTI model:

$$A_i^{\text{full}} = \text{blkdiag}(A_{R,i}, A_{P,i}, A_{Y,i}), \quad B_i^{\text{full}} = \text{blkdiag}(B_{R,i}, B_{P,i}, B_{Y,i}),$$

$$C_i^{\text{full}} = \text{blkdiag}(C_{R,i}, C_{P,i}, C_{Y,i}), \quad D_i^{\text{full}} = \text{blkdiag}(D_{R,i}, D_{P,i}, D_{Y,i}), \quad i \in \{2, 4\}.$$

In code, this is simply:

```
A_full = blkdiag(A_R, A_P, A_Y); B_full = blkdiag(B_R, B_P, B_Y);
```

```
C_full = blkdiag(C_R, C_P, C_Y); D_full = blkdiag(D_R, D_P, D_Y);
```

which yields a block-diagonal discrete-time realization $\Sigma_i^{\text{full}} = (A_i^{\text{full}}, B_i^{\text{full}}, C_i^{\text{full}}, D_i^{\text{full}})$ ready for closed-loop synthesis and frequency-domain assessment.

full model 2

$$A_2^{\text{full}} = \begin{bmatrix} 1 & 0.00099956 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.99912 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0.00099956 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.99912 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0.00099956 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.99912 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0.00099956 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.99912 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0.00099956 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.99912 \end{bmatrix}$$

$$B_2^{\text{full}} = \begin{bmatrix} 0.000178168 & 0 & 0 \\ 0.356284 & 0 & 0 \\ 0 & 3.17978e-05 & 0 \\ 0 & 0.0633581 & 0 \\ 0 & 0.000126718 & 0 \\ 0 & 0.253432 & 0 \\ 0 & 0 & 2.50939e-05 \\ 0 & 0 & 0.0501876 \\ 0 & 0 & 0.000100362 \\ 0 & 0 & 0.20075 \end{bmatrix}$$

$$C_2^{\text{full}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad D_2^{\text{full}} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

full model 4

$$A_4^{\text{full}} = \begin{bmatrix} 1 & 0.000999295 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.998591 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0.000999295 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.998591 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0.000999295 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.998591 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0.000999295 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.998591 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0.000999295 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0.998591 \end{bmatrix}$$

$$B_4^{\text{full}} = \begin{bmatrix} 0.00018233 & 0 & 0 \\ 0.364659 & 0 & 0 \\ 0 & 3.54986e-05 & 0 \\ 0 & 0.0709923 & 0 \\ 0 & 0.000141994 & 0 \\ 0 & 0.283969 & 0 \\ 0 & 0 & 2.78718e-05 \\ 0 & 0 & 0.0557324 \\ 0 & 0 & 0.000111487 \\ 0 & 0 & 0.22293 \end{bmatrix}$$

$$C_4^{\text{full}} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{bmatrix} \quad \text{and} \quad D_4^{\text{full}} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Nominal Open-Loop Step

We begin with a concise survey of the uncompensated, discrete-time step responses for the five linearized models at the synthesis points, organized by axis (Roll–Pitch–Yaw). This nominal view reveals the raw variability in static gain sign, effective time constants, and amplitude spread across the scheduled envelope. Importantly, *all plants are open-loop unstable in each axis* (each model has at least one eigenvalue outside the unit circle), so the following responses reflect divergent dynamics that will later be tamed by stabilization and augmentation. The same exercise will be repeated for the augmented models to contrast pre- and post-augmentation behavior.

Roll

The roll-axis exhibits a predominantly positive-going response with noticeable dispersion across operating points, indicating variability in loop gain and dominant time scales before feedback. Formally, *all five roll models are open-loop unstable* ($\rho(A_{R,i}) > 1$, $i = 1, \dots, 5$), which explains the unbounded trend in the nominal step.

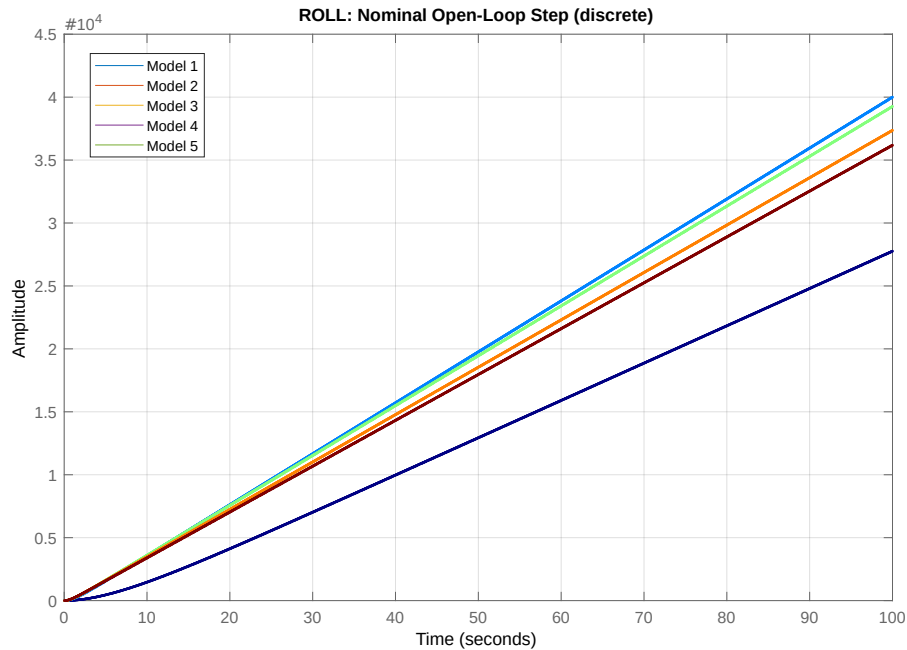


Figure 4.1 – Roll: Nominal open-loop step.

Source: Author

Downrange

Downrange responses are predominantly negative-going over a longer horizon, consistent with a negative nominal static gain, yet they also separate markedly across the five points. *All five Downrange models are open-loop unstable* ($\rho(A_{P,i}) > 1$), so the nominal time histories diverge in the absence of feedback, underscoring the need for stabilization prior to frequency-domain assessment.

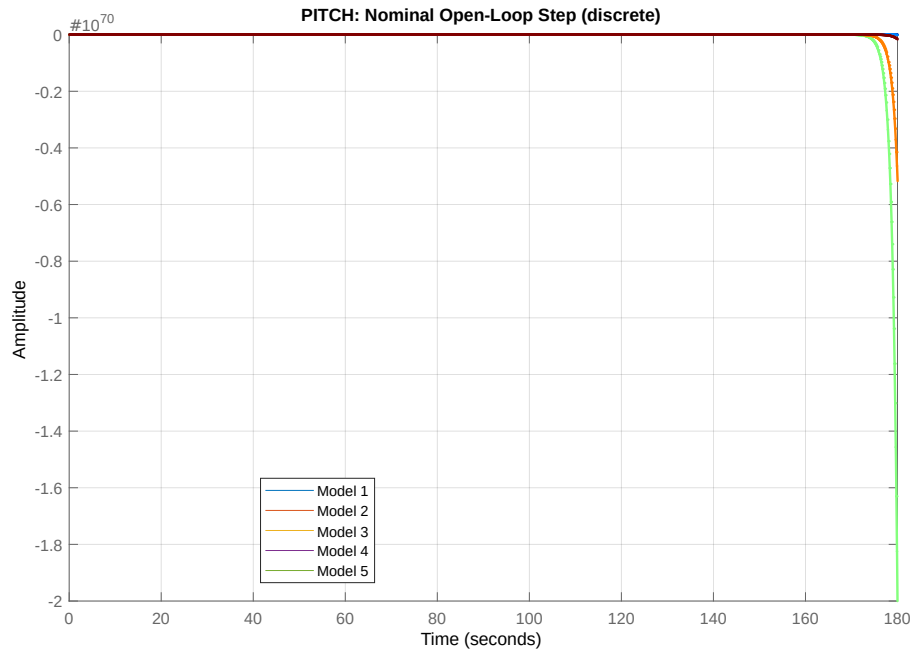


Figure 4.2 – Downrange: Nominal open-loop step.

Source: Author

Crossrange

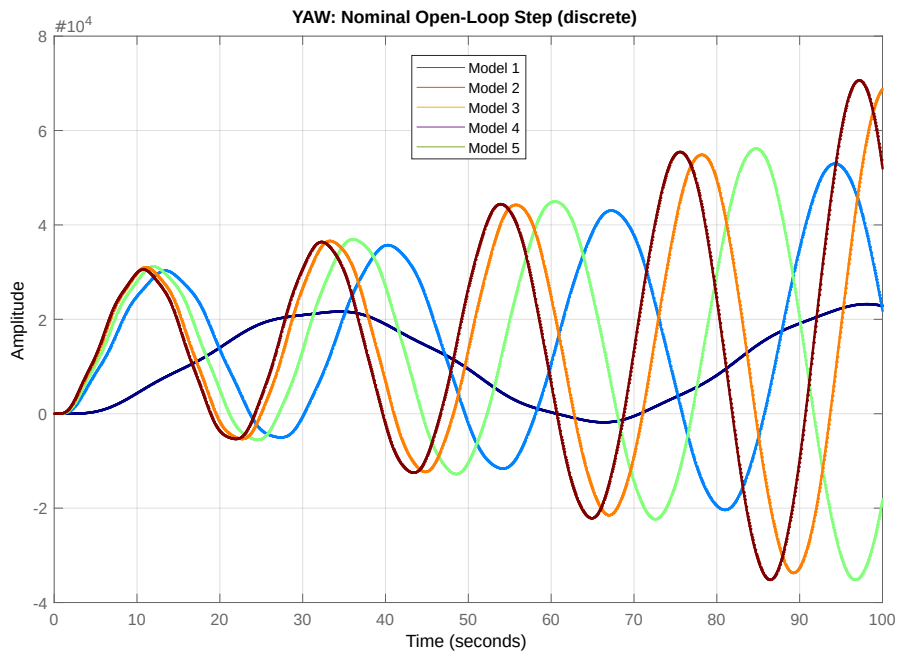


Figure 4.3 – Crossrange: Nominal open-loop step.

Source: Author

Crossrange shows the widest excursion range (both positive and negative), with clear dispersion among the five models. This axis is likewise *open-loop unstable for all five models*, suggesting lightly damped or destabilizing modes that will require careful shaping once a stabilizing controller is in place.

Augmented Model Open-Loop Step

We now examine the augmented open-loop responses obtained by appending the output-integrator/ Δu channel and the DC precompensator to each linearized plant, still without feedback. This augmentation harmonizes the steady-state scaling across operating points and exposes the shape of the slow modes that subsequent synthesis will target. As expected, augmentation alone does not confer stability: for all three axes, the discrete-time realizations remain open-loop unstable (spectral radius above unity), and long-horizon drifts are visible. Nonetheless, the overlaid traces reveal a more coherent family of step behaviors (reduced dispersion at short-to-mid horizons), which provides a clearer baseline for controller design and frequency-domain shaping to follow.

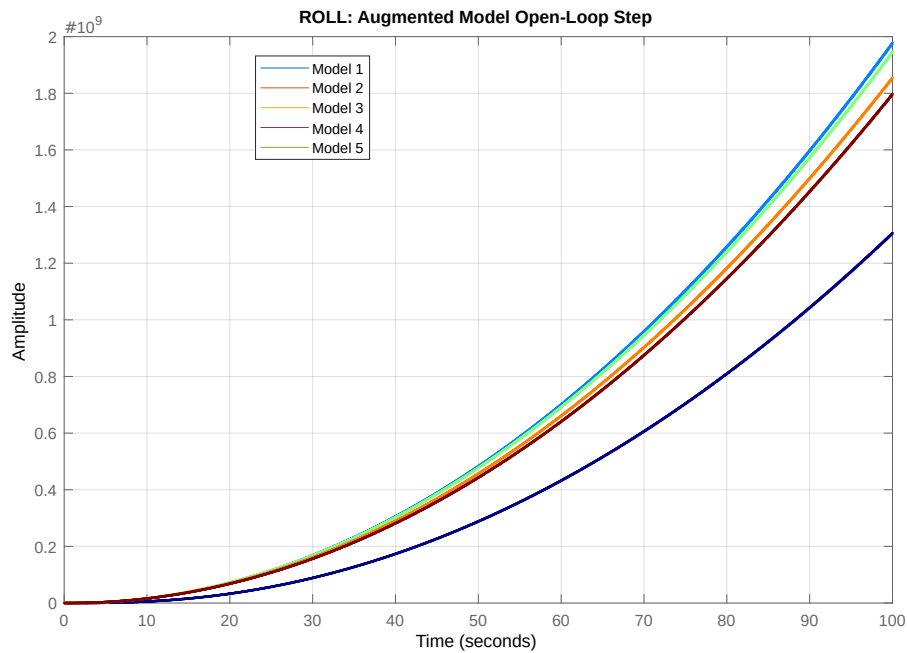


Figure 4.4 – Roll: Augmented model, open-loop step.

Source: Author

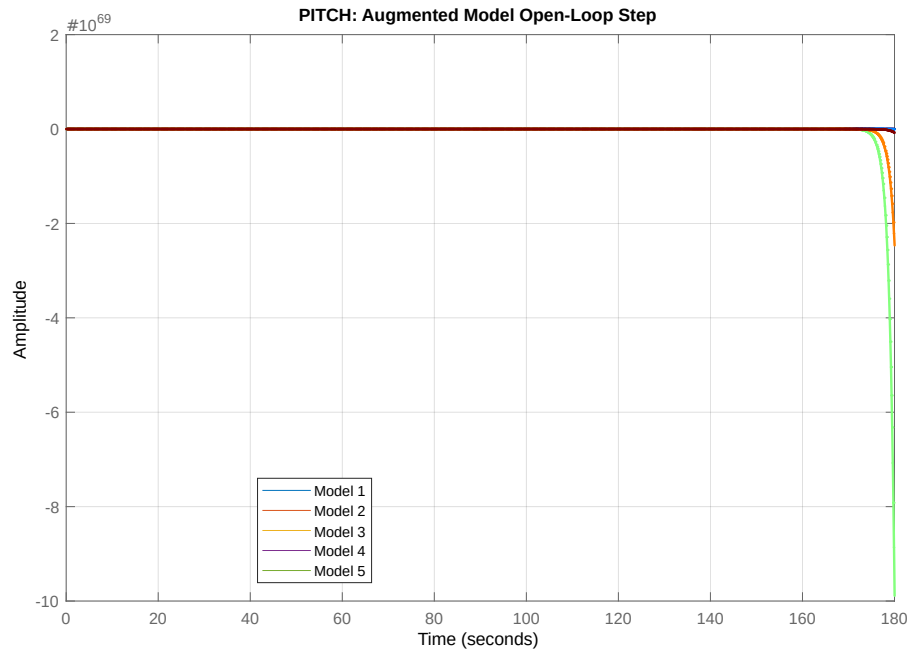


Figure 4.5 – Downrange: Augmented model, open-loop step.

Source: Author

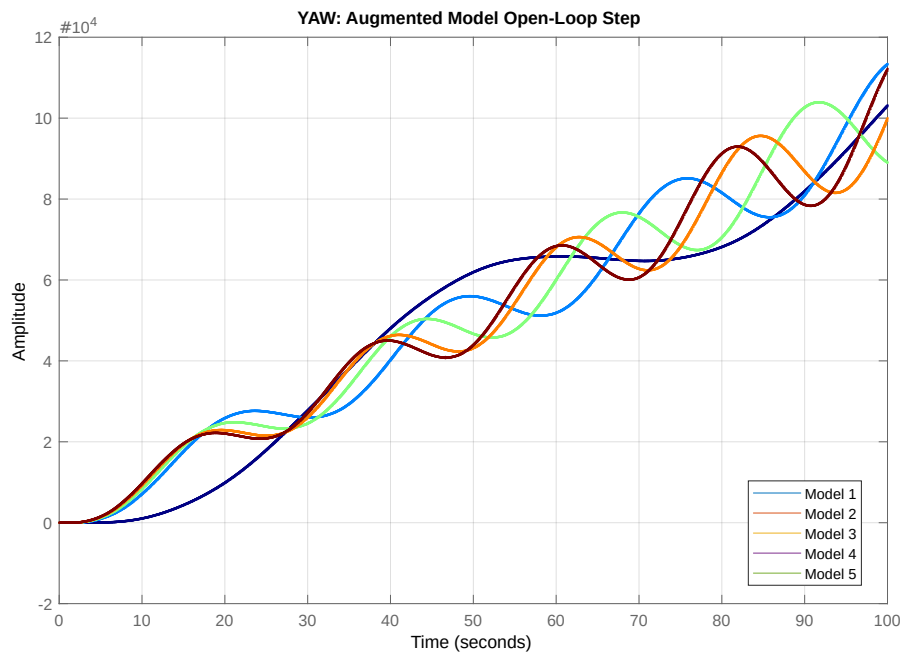


Figure 4.6 – Crossrange: Augmented model, open-loop step.

Source: Author

Controllability and Observability Check

Given the open-loop instability observed on all nominal plants, it is essential to verify that the augmented realizations at each synthesis point admit full-state regulation and estimation. The tables below summarize controllability and observability ranks per axis and model. In all cases, the ranks match the augmented state dimension n_a , confirming full controllability/observability and thus feasibility of stabilization and estimation at the analysis sampling step $T_s = 1$ ms.

Table 4.1 – Controllability/Observability ranks by axis and model (augmented realization).

Roll axis	Model	n_a	rank(Co)	rank(Ob)
	1	3	3	3
	2	3	3	3
	3	3	3	3
	4	3	3	3
	5	3	3	3
Downrange axis	Model	n_a	rank(Co)	rank(Ob)
	1	5	5	5
	2	5	5	5
	3	5	5	5
	4	5	5	5
	5	5	5	5
Crossrange axis	Model	n_a	rank(Co)	rank(Ob)
	1	5	5	5
	2	5	5	5
	3	5	5	5
	4	5	5	5
	5	5	5	5

Linear simulator (6-DoF + wind/turbulence)

We employ a linear 6-DoF simulator driven by stochastic wind inputs consistent with the Dryden and von Kármán spectral models. The mean wind profile is superposed with a zero-mean gust realization to produce altitude-dependent disturbances whose bandwidth and intensity match prescribed RMS/PSD levels. This framework supports discrete-time state-space propagation, reproducible seeds, and Monte-Carlo campaigns, and is well suited for frequency-domain validation (e.g., $|S|$, $|T|$) under small-perturbation assumptions.

The following profile, shown in Figure 4.7, summarizes the near-surface wind over Belém (AGL altitude on the vertical axis), showing the northward and eastward components, the resultant speed, and the azimuth (direction the wind blows toward) as functions of altitude.

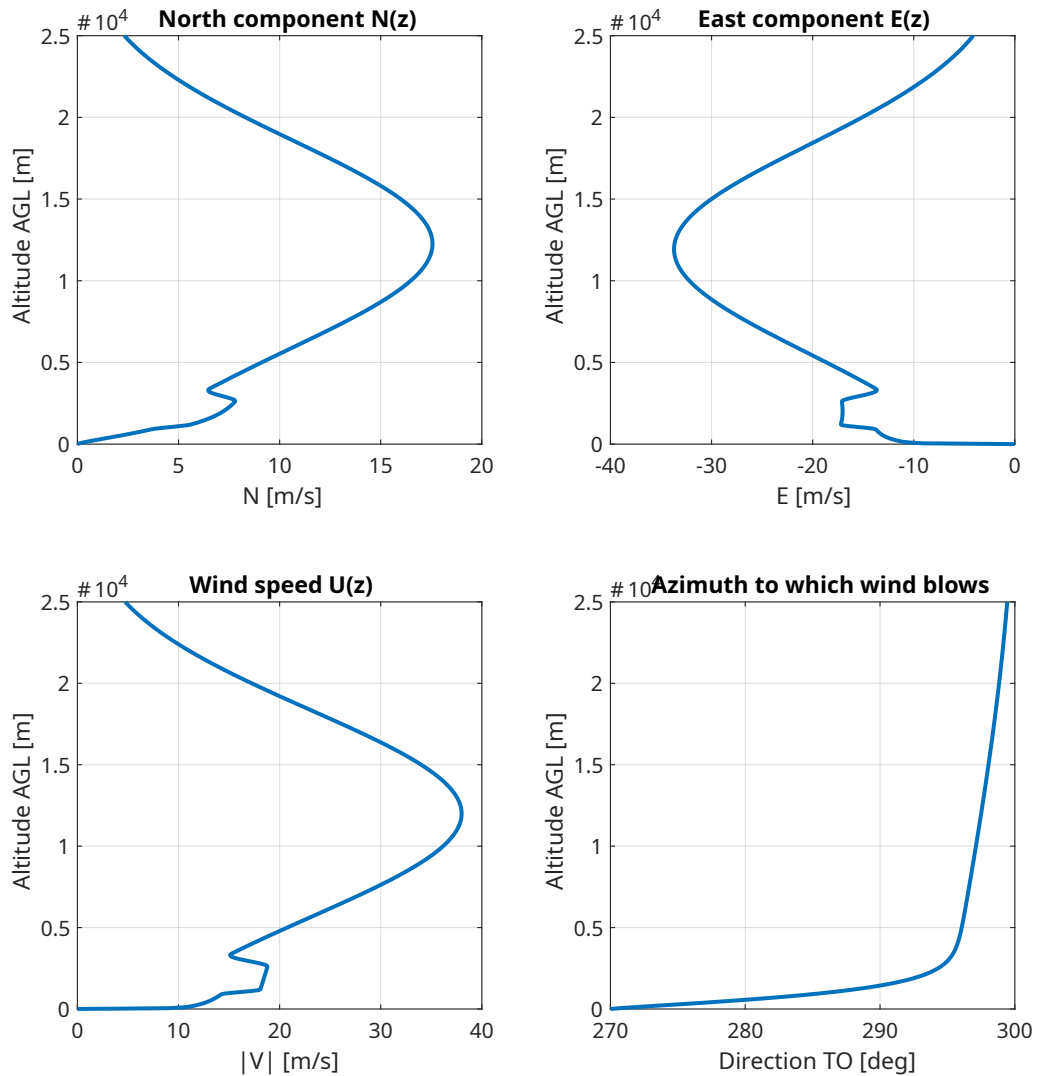


Figure 4.7 – Belém wind profile.

Source: Author

From left to right: $N(z)$ (northward component), $E(z)$ (eastward component), total wind speed $|V|(z)$, and the Direction TO azimuth. Each panel shares the same vertical axis (AGL altitude) and uses standard units (m s^{-1} for speed, degrees for direction). Together, they depict how both vector components and magnitude vary with height while the rightmost panel indicates the altitude-dependent heading toward which the wind is blowing.

A climatological mean northward wind profile was combined with a stochastic gust realization to obtain the mixed profile shown in Figure 4.8. The mean profile captures the slow, altitude-dependent trend, while the gust input injects high-frequency, zero-mean fluctuations that broaden the distribution and elevate near-surface variability.




Figure 4.8 – Mixed northward wind profile (mean + gust).

Source: Author

The solid *Mean N* trace provides the smooth background trend with height, whereas the *Mean+Gust N* curve departs locally from that trend due to injected turbulence shifting left/right around the mean and indicating stronger variability in layers where shear and turbulence intensity are higher. Together, they illustrate how superimposing a gust field on the mean profile produces a more realistic, time-instantaneous wind state for analysis and simulation

Controllers under test

We evaluate a baseline LQR design alongside LQG and LQG/LTR controllers scheduled on the operating point h (altitude). The LQR baseline provides a simple reference for bandwidth and command effort; LQG adds optimal state estimation (noise covariances tuned per axis), while LQG/LTR targets loop-transfer recovery to improve robustness and sensitivity shaping. All controllers are implemented in discrete time at $T_s = 1$ ms, with scheduling across the five synthesis points and standard safeguards (pre-filtering, rate limits, optional anti-windup). Performance is assessed via step tracking, gain/phase margins (when meaningful), and the peaks of $|S|$ and $|T|$ over the scheduled envelope.

Numerical details

All synthesis and analyses were performed in MATLAB (R2018a; Control System Toolbox, Robust Control Toolbox), using discrete-time LQR/Kalman solvers (`dlqr` and its dual), frequency-domain routines (`bode`, `sigma`, `allmargin`), and zero-order-hold

discretization (c2d('zoh')). The scheduling and closed-loop evaluations use a fixed sampling step $T_s = 1$ ms; frequency sweeps employ logarithmically spaced grids (approximately 10^{-1} – 10^2 rad/s). Experiments ran on Windows (PowerShell) on a CPU-only workstation (no GPU acceleration; standard multi-core Intel/AMD). Code corresponds to the `robust_analysis_all_mode` branch with helpers `plot_ST_colormap`, `pick_ax`, `resolve_Q/R`, and `margin_label_for_sys`; MATLAB default solver tolerances and random seeds were used unless stated otherwise.

The table 4.2 lists the flight-condition parameters that define Models 1–5; these parameters are the basis for model identification, comparison, and scheduling logic. Each model is associated with a specific operating point characterized by geometric altitude (Alt), speed (V), standard-air density (ρ), dynamic pressure ($q = \frac{1}{2}\rho V^2$), local speed of sound (a), and Mach number ($M = V/a$).

Table 4.2 – Models 1–5 and their defining flightcondition parameters.

Model	Alt [m]	V [m/s]	ρ [kg/m ³]	q [Pa]	a [m/s]	M
1	19975.67	326.13	0.089245	4746.20	295.07	1.105
2	11045.02	443.11	0.362411	35579.83	295.07	1.502
3	8744.84	435.09	0.481516	45577.03	304.89	1.427
4	6522.10	414.00	0.622731	53366.72	314.27	1.317
5	5518.48	400.67	0.696024	55869.41	318.41	1.258

4.1.2 Test Scenarios

We define a baseline descent and a set of stressed conditions to probe tracking, robustness, and actuator/sensor resilience. The nominal case follows the reference atmosphere and mass properties, with mean wind and turbulence at low intensity. Stressed cases incrementally introduce realistic gusts and actuator/sensor impairments. *Parametric uncertainty (aerodynamic derivatives, mass/inertia, density bias) is not included here and is planned for future work.* Unless otherwise stated, all cases run at the analysis sampling step T_s with identical guidance profiles and initial conditions.

- (a) **Nominal descent (reference profile).** Baseline atmosphere and density, nominal mass/inertia, mean wind only; serves as the reference trajectory for bandwidth/effort comparison.
- (b) **Gusts and shear: light / moderate / severe.** Stochastic turbulence (Dryden or von Kármán) superposed on the mean profile with increasing RMS/PSD levels; include vertical/lateral shear segments to test sensitivity peaking and disturbance rejection.

- (c) **Actuator stress and sensor imperfections.** Command rate limits and hard saturation; add pure delays where applicable. Inject sensor bias and Gaussian noise consistent with hardware specification.

4.1.3 Performance Indices

We report tracking, effort, loads, and robustness using consistent discrete-time definitions (sampling step T_s) and their continuous-time counterparts when needed.

- **Tracking error (RMS, IAE).** For a reference r_k and output y_k , $e_k = r_k - y_k$. RMS error: $\text{RMS}(e) = \sqrt{\frac{1}{N} \sum_{k=1}^N e_k^\top e_k}$.
Integral of absolute error (IAE): $\text{IAE}(e) = \sum_{k=1}^N \|e_k\|_1 T_s$.
- **Control effort.** L_2 norm (energy): $\|u\|_2 = \sqrt{\sum_{k=1}^N u_k^\top u_k T_s}$. Total deflection (absolute travel): $\text{TV}(u) = \sum_{k=2}^N \|u_k - u_{k-1}\|_1$. Saturation time: fraction of samples with any channel at its limit, $\tau_{\text{sat}} = \frac{1}{N} \sum_{k=1}^N \mathbf{1}\{|u_{k,i}| \geq u_i^{\text{max}} \text{ for some } i\}$.
- **Robustness.** Classical gain/phase margins from loop transfer via `allmargin`. Sensitivity and complementary sensitivity peaks: $\|S\|_\infty = \sup_{\omega} \bar{\sigma}(S(e^{j\omega T_s}))$, $\|T\|_\infty = \sup_{\omega} \bar{\sigma}(T(e^{j\omega T_s}))$ (max singular value).

4.2 Robustness and Frequency-Domain Analysis

4.2.1 LQR Loop Shapes and Margins

The robustness properties of the scheduled LQR controllers were assessed by analyzing the closed-loop shapes, sensitivity $S(e^{j\omega T_s})$ and complementary sensitivity $T(e^{j\omega T_s})$ functions, as well as the associated gain and phase margins at representative points along the trajectory. The evaluation covers the Downrange, roll, and Crossrange axes separately, with comparisons across five linearized models per axis.

Step Responses

Figures 4.9–4.11 present the closed-loop step responses for Downrange, roll, and Crossrange, respectively. The controllers provide a satisfactory transient behavior across all scheduled models. The Downrange and Crossrange channels show moderate rise times

with bounded overshoot, while the roll channel demonstrates a faster response, as expected from the higher natural frequency of lateral dynamics.




Figure 4.9 – Step responses for the Downrange LQR controllers (five models).

Source: Author




Figure 4.10 – Step responses for the scheduled Roll LQR controllers.

Source: Author




Figure 4.11 – Step responses for Crossrange LQR controllers.

Source: Author

Singular Value Analysis






Figure 4.12 – Singular values of T and S for Downrange.

Source: Author

Figure 4.13 – Singular values of T and S for Roll.

Source: Author

Figure 4.14 – Singular values of T and S for Crossrange.

Source: Author

The singular values of $T(e^{j\omega T_s})$ and $S(e^{j\omega T_s})$ are reported in Figures 4.12–4.14. The complementary sensitivity T curves remain below 0 dB at high frequency, limiting

measurement noise amplification, while the sensitivity S exhibits low magnitude in the low-frequency range, indicating effective disturbance rejection. Across all axes, the spread between models is limited, confirming robustness of the scheduled controller set.

Stability Margins

Closed-loop Bode diagrams are shown in Figures 4.15–4.17. The corresponding numerical gain and phase margins are summarized in Table 4.3. The results highlight that:

- Downrange margins are consistent across models, with GM \approx 6 dB and PM ranging between 40.8° and 51.9° .
- Roll controllers yield a uniform margin set, GM \approx 6 dB and PM \approx 46° across all models.
- Crossrange margins are slightly lower, with PM decreasing from 34.4° to 34.1° , though still acceptable in practice.




Figure 4.15 – Closed-loop Bode diagram with margins for Downrange.

Source: Author




Figure 4.16 – Closed-loop Bode diagram with margins for roll.

Source: Author

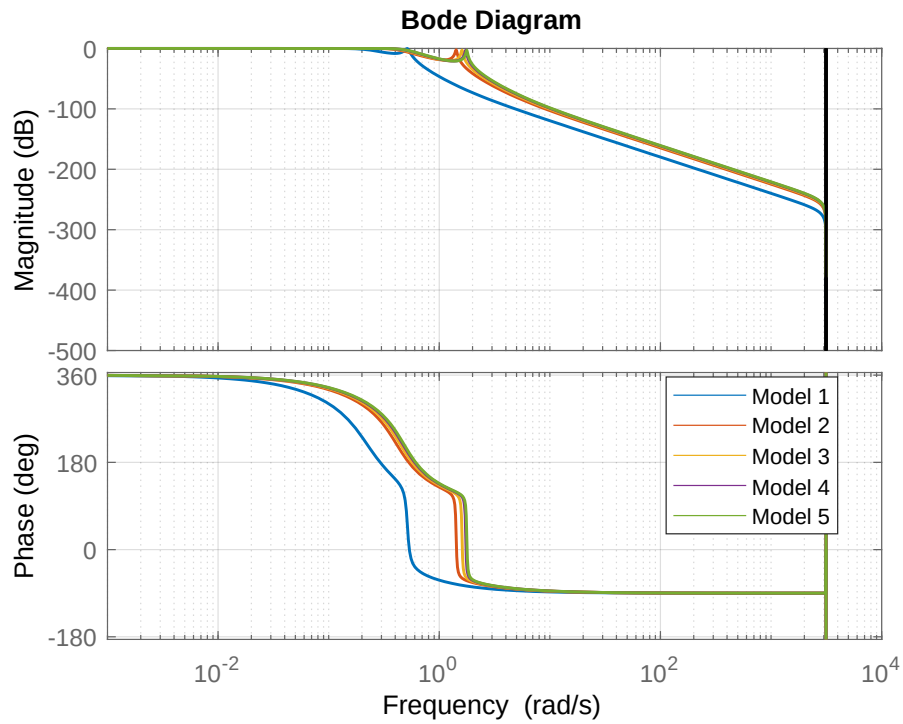


Figure 4.17 – Closed-loop Bode diagram with margins for Crossrange.

Source: Author

Table 4.3 – LQR gain and phase margins at scheduled points.

Axis	GM [dB]	PM [deg] (range)	Remarks
Downrange	6.02	40.8 – 51.9	Robust margins, consistent
Roll	6.02	45.9 – 46.1	Uniform across models
Crossrange	6.02	34.1 – 34.4	Lower phase margin

Overall, the scheduled LQR controllers guarantee acceptable robustness margins for all three axes. The roll and Downrange axes display consistent stability reserves across the scheduled models. The Crossrange axis is comparatively more critical, with reduced phase margin. In particular, the boost-back maneuver drives the vehicle predominantly along the downrange direction; consequently, the available aerodynamic authority for steering downrange exceeds that for crossrange corrections. In the crossrange channel, lateral velocities are comparatively small, yielding insufficient dynamic pressure to achieve substantial repositioning without risking oscillatory behavior. Hence, although the crossrange axis is more critical, the primary control objective there is regulation, i.e., disturbance rejection and stabilization, rather than large crossrange repositioning.

4.2.2 Kalman Filter Loop Shapes and Stability Margins

The robustness characteristics of the Kalman Filter (KF) augmented loops were assessed by analyzing the Closed-loop shapes, sensitivity $S(e^{j\omega T_s})$ and complementary sensitivity $T(e^{j\omega T_s})$, as well as the gain and phase margins at scheduled operating points. Results are reported for the Downrange, roll, and Crossrange channels separately.

Step Responses

Figures 4.18–4.20 show the step tracking responses under KF estimators. All three channels exhibit fast convergence with limited overshoot. The Downrange and Crossrange estimators provide smooth transients with rise times below one second, whereas the roll estimator achieves a sharper response due to the faster lateral dynamics.

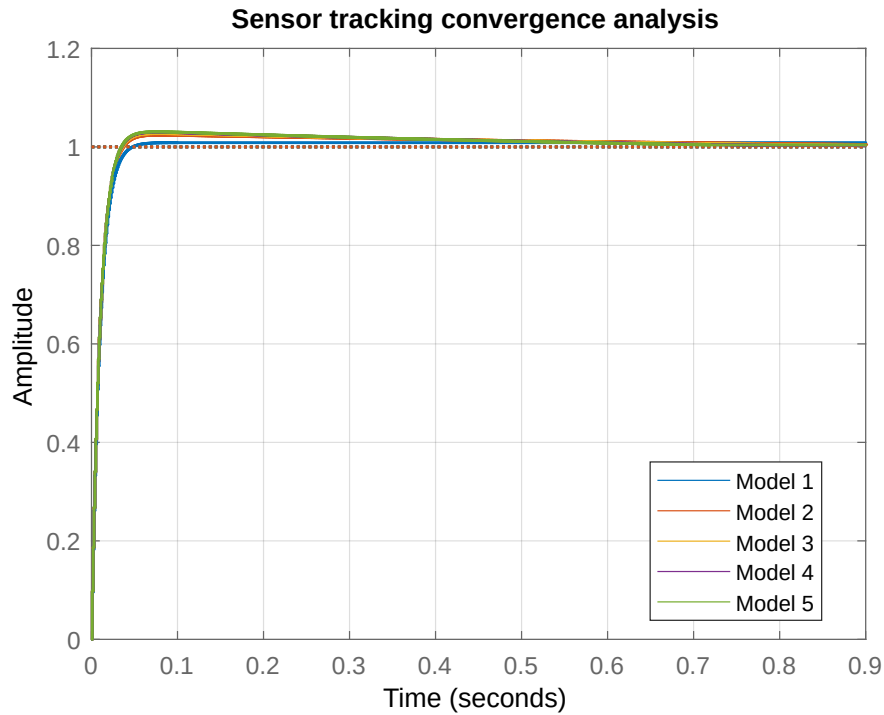


Figure 4.18 – Step responses for the scheduled Downrange KF estimators.

Source: Author




Figure 4.19 – Step responses for the scheduled roll KF estimators.

Source: Author




Figure 4.20 – Step responses for the scheduled Crossrange KF estimators.

Source: Author

Singular Value Analysis




Figure 4.21 – Singular values of T and S for Downrange.

Source: Author




Figure 4.22 – Singular values of T and S for roll.

Source: Author




Figure 4.23 – Singular values of T and S for Crossrange.

Source: Author

The singular values of $T(e^{j\omega T_s})$ and $S(e^{j\omega T_s})$ are reported in Figures 4.21–4.23. In all cases, T remains well below 0 dB at high frequencies, preventing excessive noise amplification, while S demonstrates effective disturbance rejection in the low-frequency range. Model-to-model variations remain limited, suggesting good robustness to scheduling.

Stability Margins.

Closed-loop Bode diagrams are shown in Figures 4.24–4.26. The corresponding numerical margins are summarized in Table 4.4.

Table 4.4 – KF estimators: gain and phase margins across scheduled models.

Axis	GM [dB]	PM [deg] (range)	Remarks
Downrange	5.65–5.91	54.5–58.4	Robust, consistent margins
Roll	4.92	44.7–44.8	Slightly reduced gain margin
Crossrange	6.01–6.02	56.9	High phase margin




Figure 4.24 – Closed-loop Bode diagram with margins for Downrange KF estimators.

Source: Author




Figure 4.25 – Closed-loop Bode diagram with margins for roll KF estimators.

Source: Author




Figure 4.26 – Closed-loop Bode diagram with margins for Crossrange KF estimators.

Source: Author

The KF-based loop shaping improves phase margins significantly compared to standard LQR designs, particularly in the Downrange and Crossrange channels, which both exceed 55° . The roll channel remains the most limiting case, with margins slightly below those of LQR. Nevertheless, the results confirm that the KF augmentation provides a satisfactory balance between robustness and tracking accuracy, making it a strong candidate for integration into the complete LQG/LTR framework.

4.2.3 LQG/LTR Loop Shapes and Stability Margins

This section presents the robustness assessment of the scheduled LQG/LTR controllers per axis. The closed-loop shapes was reported via Bode diagrams, the sensitivity $S(e^{j\omega T_s})$ and complementary sensitivity $T(e^{j\omega T_s})$ through singular-value plots, and the gain/phase margins evaluated at representative scheduled operating points.

Step Responses

Figures 4.27–4.29 show the closed-loop step responses obtained at five scheduled models for each axis. The Downrange and roll channels exhibit consistent transients with bounded overshoot and short settling times; the Crossrange axis displays slower dynamics in line with its lower phase reserve.




Figure 4.27 – Step responses for scheduled Downrange LQG/LTR controllers.

Source: Author




Figure 4.28 – Step responses for scheduled roll LQG/LTR controllers.

Source: Author




Figure 4.29 – Step responses for scheduled Crossrange LQG/LTR controllers.

Source: Author

Singular-Value Analysis of S and T

Figures 4.30–4.32 report the singular values of $T(e^{j\omega T_s})$ (noise/command amplification) and $S(e^{j\omega T_s})$ (disturbance rejection). Across axes, S remains low at low frequency, indicating effective disturbance attenuation, while T stays below 0 dB at high frequency, limiting sensor-noise amplification. The inter-model spread is modest, confirming robustness across the gain-scheduled set.






Figure 4.30 – Singular values of T and S for Downrange.

Source: Author

Figure 4.31 – Singular values of T and S for roll.

Source: Author

Figure 4.32 – Singular values of T and S for Crossrange.

Source: Author

Closed-Loop Shapes and Margins




Figure 4.33 – Closed-loop Bode diagram and margins for Downrange LQG/LTR.

Source: Author




Figure 4.34 – Closed-loop Bode diagram and margins for roll LQG/LTR.

Source: Author




Figure 4.35 – Closed-loop Bode diagram and margins for Crossrange LQG/LTR.

Source: Author

Bode diagrams with computed stability margins are shown in Figures 4.33–4.35. The numerical gain and phase margins at the scheduled points are summarized in Table 4.5.

Table 4.5 – LQG/LTR gain and phase margins across scheduled models.

Axis	GM [dB]	PM [deg] (range)	Remarks
Downrange	6.02	40.77–51.93	Consistent, satisfactory reserves
Roll	6.02	45.80–46.01	Uniform across models
Crossrange	6.02	34.10–34.37	Lowest PM; most critical axis

Overall, the scheduled LQG/LTR controllers ensure acceptable robustness for Downrange and roll, with GM \approx 6 dB and phase margins in excess of 45° (roll) and up to \sim 52° (Downrange). The Crossrange channel is the most constraining case, with PM \approx 34°, which is still workable but leaves less attenuation to unmodelled dynamics and delays. If higher Crossrange robustness is required, loop-shape refinement (e.g., estimator noise covariances or recovery weighting) can be used to trade a small bandwidth reduction for a larger phase reserve.

4.2.4 LTR Turning Effect

The LTR procedure was applied via estimator shaping to pull the open-loop seen by the controller toward the LQR target around the crossover, with the expected

impact of recovering slope and phase. The data confirm this effect most clearly on the Crossrange and Downrange channels. Compared to the baseline LQR margins (Pitch: $\text{GM} \approx 6.02$ dB, $\text{PM} = 40.851.9^\circ$; Yaw: $\text{GM} \approx 6.02$ dB, $\text{PM} = 34.134.4^\circ$), the KF-shaped loops (our practical LTR turning) increase the phase reserve to $\text{PM} = 54.558.4^\circ$ in Downrange and to $\sim 56.9^\circ$ in Crossrange, while preserving gain margins close to 6 dB on Crossrange and near 5.75.9 dB on Downrange (see Table 4.4, and the margin plots in Figs. 4.24–4.26). This aligns with the singular-value trends: $S(e^{j\omega T_s})$ remains low at low frequency (disturbance rejection preserved), and $T(e^{j\omega T_s})$ rises moderately around crossover as the loop opens to recover phase (Figs. 4.21–4.23). On roll, the KF/LTR effect is more neutral: the phase margin stays near 45° and the gain margin slightly reduces to ~ 4.92 dB, indicating that aggressive recovery offers limited benefit on this axis relative to its LQR/LQG shape (see Tables 4.3–4.5).

Overall, the evidence supports the canonical LTR trade-offs. Where the baseline LQR phase was tight (notably Crossrange), LTR turning yields a substantial phase gain (from $\sim 34^\circ$ to $\sim 57^\circ$) with loop shapes that remain well-behaved across the scheduled models. The price is the expected increase of $T(e^{j\omega T_s})$ at mid/high frequency (noise passage and control effort), which should be managed by partial recovery and bandwidth moderation rather than full recovery. In practice, the present settings achieve the intended recovery of the LQR-like loop around ω_c for Downrange and Crossrange, while keeping S/T spreads small between models and maintaining acceptable stability reserves across the schedule (cf. Figs. 4.12–4.14 vs. Figs. 4.21–4.23).

4.3 Trajectory-Wide Cross-Regime Comparison

Given that the roll axis is intentionally constrained at 2π rad during the aerodynamic steering and the Crossrange shares the same control synthesis and scheduling structure as the Downrange axis (time-based scheduling, LQG/LTR observer-controller, and on-switch K_{comp} adaptation), we adopt the it, Downrange, as the *representative case* throughout. Unless stated otherwise, methods, weights (Q/R), sampling T_s , and implementation details are identical for them; only the calibrated state-space matrices differ. Accordingly, all analyses, figures, and conclusions are presented for *Downrange* and carry over to *roll Mutatis mutandis*¹; any axis-specific effects will be noted when relevant.

Figures 4.36 and 4.37 overlay the closed-loop responses obtained by freezing each scheduled model and simulating it over the entire descent profile under identical reference commands without disturbances. This all-model test is intentionally conservative: each controller is exercised outside its nominal synthesis region, so differences in loop gain and bandwidth become explicit.

¹ (Lat.): with the necessary changes having been made.

1. **Sensitivity of Model 1 at major manoeuvres.** Model 1 exhibits the largest output excursions and control activity at the sharpest reference waypoints. This is consistent with its synthesis point being in the supersonic regime ($M > 1$), where (i) aerodynamic effectiveness is higher and (ii) the selected LQ weights typically deliver a higher loop crossover for rapid disturbance rejection. The resulting complementary sensitivity $T(z)$ peaks closer to the target bandwidth, which improves tracking in its design bin but appears as over-reactivity when the plant drifts away from that bin.
2. **Smoother behaviour of the other models.** Controllers synthesized in the transonic/subsonic bins show more gradual responses and smaller control deflections at the same waypoints. Their lower effective loop gain outside their design points reduces sensitivity and helps avoid excessive actuator activity or rate saturation when dynamic pressure is lower.
3. **Steady-state alignment.** When the “Weighted” variant is used, the static compensation factor K_{comp} eliminates the DC bias, so residual differences across models are dynamic (bandwidth/phase) rather than static (gain).
4. **State evolution.** The state plots confirm that the integral-augmented estimator keeps the reconstructed states bounded across bins; overshoots correlate with the same manoeuvre points highlighted in the output plot, again with Model 1 showing the sharpest transients.




Figure 4.36 – All-model overlay: reference tracking and control effort across the full trajectory.

Source: Author

The observed dispersion is not a flaw but the expected signature justifying the scheduling: closed-loop properties are optimal in the local synthesis region and naturally deviate from it. In particular, a supersonic design (Model 1) favors a rapid decay of the error, useful for preventing large aerodynamic excursions at high dynamic pressure, but this same aggressiveness may appear "sensitive" when applied to subsonic conditions.

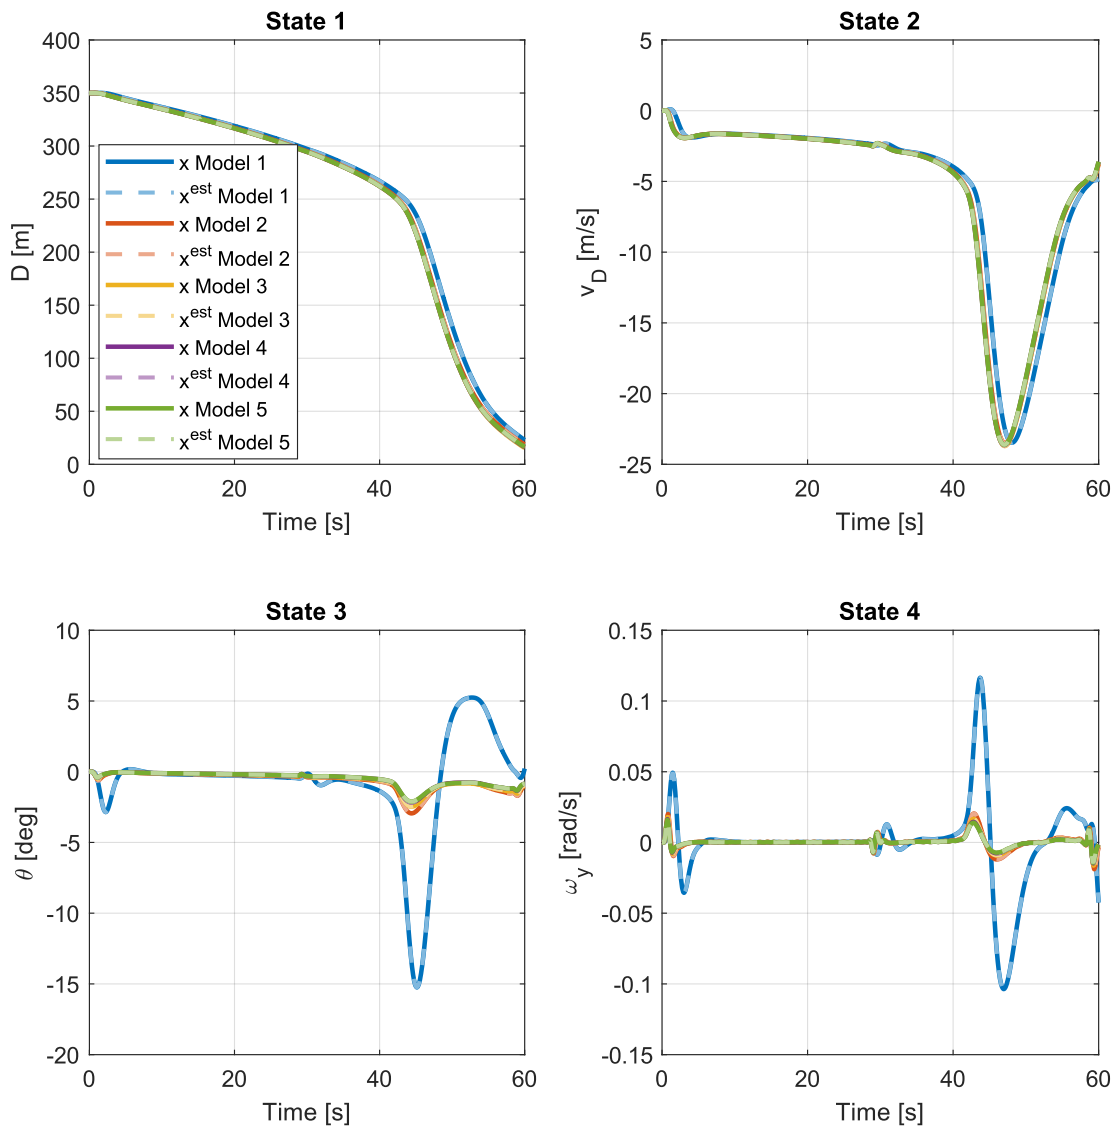


Figure 4.37 – All-model overlay: state evolution (pairs) over the same run.

Source: Author

4.4 Scheduling and Implementation Aspects

4.4.1 Mapping and selection rule.

Let the scheduled synthesis instants be $\{t_j\}_{j=1}^{N_{sc}}$ (stored in $\mathbf{S.t}$, in seconds), and let the descent profile be $h(t)$ (altitude as a function of time). We assign each instant t of the trajectory and by composition each altitude $h(t)$ to a unique scheduled model index j by a time-based partition with mid-point boundaries:

$$\mathcal{I}_1 = \left[0, \frac{t_1+t_2}{2} \right), \quad \mathcal{I}_j = \left[\frac{t_{j-1}+t_j}{2}, \frac{t_j+t_{j+1}}{2} \right), \quad j = 2, \dots, N_{sc} - 1, \quad \mathcal{I}_{N_{sc}} = \left[\frac{t_{N_{sc}-1}+t_{N_{sc}}}{2}, t_{\text{final}} \right].$$

At any time $t \in \mathcal{I}_j$ (equivalently at altitude $h(t)$), the controller uses model j . This deterministic time-to-index mapping avoids chatter and directly induces an altitude-to-index map via $h(t)$.

Scheduling points are denser where the plant varies most rapidly with flight condition classically near the *transonic corridor* ($M \approx 0.81.2$) and around *high dynamic pressure* $q = \frac{1}{2}\rho V^2$. In those regions aerodynamic derivatives and control effectiveness exhibit strong gradients; increasing the local density of operating points reduces inter-model mismatch and prevents excess gain/phase dispersion across the envelope.




Figure 4.38 – Operating-point coverage along the descent: active model index by time (hence altitude), with synthesis points (black dots) and switching boundaries (light grey dashed lines).

Source: Author

Figure 4.38 overlays the descent profile with the scheduled coverage: (i) colored segments indicate which model index j is active over each interval \mathcal{I}_j ; (ii) small black circles

mark the synthesis instants $\{t_j\}$; (iii) light grey dashed vertical lines show the mid-point boundaries $\partial\mathcal{L}_j$ where the active model switches. Two observations follow. First, the scheduled points cluster around the transonic / high- q portions of the trajectory precisely where $h(t)$ and $V(t)$ induce the fastest variations of aero-damping and control authority yielding short intervals and finer local fidelity. Second, the transitions align with physically meaningful altitude bands; the output tracks the reference without discontinuities at the boundaries, indicating that the augmented integrator and the DC precompensation keep inter-segment biases small. Residual offset, if any, is predominantly due to inter-model mismatch and is further mitigated by the (optional) K_{comp} adaptation applied at the boundaries.

4.4.2 Reference-Tracking Refine Strategies under Time-Scheduled Control

Each panel (Figs. 4.39-4.41) overlays the reference $r(t)$ with the measured output $y(t)$ (solid) and its internal estimate $\hat{y}(t)$ (dashed). Colored segments indicate the active scheduled model; small black circles mark the synthesis instants t_j ; light-grey dashed lines show the mid-point switching boundaries between models (time-based mapping via $S.t$).

Nominal (baseline).




Figure 4.39 – Reference $r(t)$, measured output $y(t)$ (solid), and internal estimate $\hat{y}(t)$ (dashed) Nominal.

Source: Author

Without look-ahead or boundary reweighting, the nominal controller (Fig. 4.39) exhibits a modest rise time and conservative overshoot. Across model transitions, a small steady-state offset accumulates (bias relative to r) but remains bounded; switching is smooth thanks to the augmented integrator and DC pre-compensation. This plot serves as the baseline tracking envelope.

Predictive (look-ahead).

With a d -step look-ahead on the reference, the response (Fig. 4.40) anticipates set-point trends and disturbance entries. The output aligns earlier with $r(t)$ around sharp features, reducing lag and integral error within each interval. As expected, this sharper tracking comes with visibly stronger control activity (you will see larger excursions in the control subplot), though transitions at boundaries remain continuous.




Figure 4.40 – Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ Predictive.

Source: Author

Weighted (boundary compensation).

Applying the learned/static K_{comp} on the servo path (Fig. 4.41) mitigates inter-model biases. The output adheres closely to the reference after each switch, flattening residual offset while preserving the nominal smoothness in-interval. Compared to the baseline, the steady-state error is notably reduced around regime changes, with only a mild increase in command effort.

Within each scheduled interval, *Predictive* yields the tightest tracking to $r(t)$; *Weighted* systematically improves post-switch alignment; *Nominal* is most conservative and shows the largest residual bias near boundaries. All three maintain continuity at switches, indicating that the augmented observer and pre-compensation are effective at preventing discontinuities.




Figure 4.41 – Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ Weighted.

Source: Author

These impressions were quantified with standard metrics J_{LQR} , IAE/ISE/ITAE/ITSE, RMSE, and effort terms $\int u^2 dt$, $\int (\Delta u)^2 dt$ to compare tracking quality and control activity across the Nominal, Predictive, and Weighted strategies.

4.4.3 Performance

We evaluate three control variants—*Nominal*, *Predictive*, and *Weighted*—over the full trajectory. Integral tracking indices (IAE/ISE/ITAE/ITSE), the discrete-time LQR cost $J_{\text{LQR}} = \sum_k (x_a^\top Q x_a + \Delta u^\top R \Delta u) T_s$, and effort/steadiness metrics are reported below.

Table 4.6 – Optimal cost values J_{LQR} for control strategies

Method	J_{LQR}
Weighted	4.3380e+03
Nominal	4.1756e+03
Predictive	6.4302e+03

Table 4.7 – Performance without disturbance.

(a) Tracking metrics.

Method	IAE	ISE	ITAE	ITSE	RMSE
Weighted	1.4671×10^2	5.6036×10^2	5.0901×10^3	2.2952×10^4	3.0560
Nominal	2.5225×10^2	1.7306×10^3	9.5518×10^3	7.6660×10^4	5.3706
Predictive	9.6176×10^1	2.0140×10^2	2.9614×10^3	6.9911×10^3	1.8321

(b) Control effort and steady-state figures.

Method	$\sum_{k=1}^N u_k^2 T_s$	$\sum_{k=1}^N (\Delta u_k)^2 T_s$	Steady-state error (mean \pm std)	$ u _{\max}$
Weighted	2.2656×10^{-1}	3.1068×10^1	-2.1368 ± 1.5169	0.326
Nominal	2.2058×10^{-1}	2.9801×10^1	-4.8219 ± 2.2124	0.317
Predictive	2.4754×10^{-1}	5.1803×10^1	-2.0033 ± 1.7841	0.378

Three salient points emerge:

1. **Predictive** delivers the *best tracking* (lowest IAE/ISE/ITAE/ITSE and RMSE = 1.83), but at the price of a *higher actuation activity*: its J_{LQR} and $\int (\Delta u)^2 dt$ are the largest (consistent with the look-ahead action increasing responsiveness).
2. **Weighted** improves tracking markedly over *Nominal* (e.g., RMSE 3.06 vs. 5.37; ISE reduced by $\sim 68\%$) while keeping the actuation energy at a moderate level and maintaining the peak command within a similar range.
3. **Nominal** exhibits the *largest steady-state bias* (mean error ≈ -4.82) and the weakest tracking, although its control effort remains modest.

If the design objective prioritizes tracking accuracy, *Predictive* is preferable; if a balanced trade-off between tracking and control activity is sought, *Weighted* offers the most attractive compromise. The *Nominal* baseline is conservative but underperforms on accuracy and steady-state residuals.

4.5 Stress Testing and Sensitivity

Across all three configurations (Nominal, Predictive, and Weighted), the LQG/LTR controller remains robust throughout the descent: no actuator saturation is observed, no runaway behavior occurs, and the control effort stays bounded. A small oscillation appears in the transonic region immediately after the first model switch. This is consistent with inter-model mismatch at a rapidly varying operating point and with the deliberate choice, in this dissertation, to *avoid* gain interpolation between scheduled models.




Figure 4.42 – Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ under injected disturbances: Nominal.

Source: Author




Figure 4.43 – Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ under injected disturbances: Predictive.

Source: Author

Importantly, the oscillation is *bounded* and quickly damped, indicating preserved phase and gain margins. Rather than a deficiency, it is an informative feature: it reveals




Figure 4.44 – Reference $r(t)$, measured output $y(t)$, and internal estimate $\hat{y}(t)$ under injected disturbances: Weighted.

Source: Author

where transition smoothing would provide the most benefit without compromising the core robustness achieved here. As such, it motivates a natural line of future work:

- introduce smooth gain blending across the schedule (e.g., time/altitude-based interpolation of K , L , K_{servo} , P_c , and/or K_{comp});
- apply bumpless-transfer refinements (to remove any small bias/step at boundaries);
- tune dwell time and hysteresis at regime boundaries to further suppress boundary-induced ripples;
- optionally adjust the K_{comp} adaptation (filtering and step response) specifically near the transonic/high- q corridor.

Overall, the figures and metrics confirm effective rejection of wind disturbances in all cases, while also indicating clear opportunities to refine transition behavior in the most demanding parts of the envelope.

Table 4.8 – Optimal cost values J_{LQR} for control strategies under disturbance

Method	J_{LQR}
Weighted	7.7695e+03
Nominal	9.5908e+03
Predictive	1.0313e+04

Table 4.9 – Performance under disturbance.

(a) Tracking metrics.

Method	IAE	ISE	ITAE	ITSE	RMSE
Nominal	4.3688×10^2	9.6620×10^3	1.6338×10^4	3.6880×10^5	12.690
Predictive	2.6094×10^2	5.1449×10^3	9.1146×10^3	1.8905×10^5	9.260
Weighted	2.8618×10^2	4.7218×10^3	1.0248×10^4	1.7632×10^5	8.871

(b) Control effort and steady-state figures.

Method	$\sum_{k=1}^N u_k^2 T_s$	$\sum_{k=1}^N (\Delta u_k)^2 T_s$	SS mean	SS std	$ u _{\max}$
Nominal	3.7889×10^{-1}	6.5849×10^1	-4.2517	1.4382	0.378
Predictive	3.9301×10^{-1}	8.3509×10^1	-1.4111	1.1306	0.383
Weighted	4.3221×10^{-1}	9.0919×10^1	-2.0443	1.3253	0.369

- **Tracking under disturbance.** *Predictive* achieves the best integral tracking (lowest IAE/ITAE/ITSE), closely followed by *Weighted* on squared error: it actually attains the lowest ISE and the best RMSE in Table 4.9. *Nominal* remains most conservative and shows the largest residual bias.
- **Control activity.** Disturbance rejection increases actuation. Both *Predictive* and *Weighted* show higher $\sum (\Delta u_k)^2 T_s$ than *Nominal*, with *Weighted* the largest, consistent with its post-switch bias flattening via K_{comp} . Peak command $|u|_{\max}$ stays comparable across strategies.
- **Cost balance.** J_{LQR} grows from *Nominal* to *Predictive* to *Weighted*, reflecting the trade between sharper tracking (especially at boundaries) and increased control effort/augmented-state excursion.

4.6 Summary

Under injected disturbances, the *Predictive* strategy minimizes integral tracking indices, while the *Weighted* strategy yields the smallest squared tracking error (ISE) and RMSE at the cost of the highest control activity. The *Nominal* controller serves as a conservative baseline with the largest residual error but the lowest overall activity.

5 Conclusions

5.1 Main Achievement

This work set out to design, implement, and critically assess a practical LQG/LTR-based control framework for the aerodynamic descent of a reusable launch vehicle (RLV), with explicit support for time-scheduled modeling along a representative trajectory. Beyond the synthesis itself, emphasis was placed on a reproducible simulation pipeline, transparent diagnostics, and performance assessment under realistic disturbances. The main achievements can be summarized as follows.

A coherent, R2018compatible simulation stack.

We developed the `SIM_v02` pipeline, a compact MATLAB implementation that (i) runs parallel simulations for each axis/model, (ii) remains fully compatible with MATLAB R2018a (no `tiledlayout` dependency, careful plotting fallbacks), and (iii) offers a uniform interface for three controller variants: *Nominal*, *Predictive* (lookahead), and *Weighted* (boundary compensation). A consistent plotting style was enforced so that `SIM_v02` reproduces the figure composition of the original `code_souche` (two figures: output/control and states). Titles, axis units (including angles in degrees), and legends were standardized for readability across axes and models.

Augmented LQG/LTR design with DC precompensation.

For each scheduled model, we constructed the same augmented architecture: an integrator on the controlled output, a discrete-time LQR law acting on augmented states, and a dual KF observer. We explicitly derived the static DC pre-compensator $P_c = 1/K_s$ from the plants steadystate gain $K_s = C_d(I - A_d)^{-1}B_d + D_d$, ensuring bumpless incremental actuation and unbiased steady outputs in nominal conditions. The implementation includes guards for ill-conditioned cases and normalizes D_d shapes to avoid dimension mismatches.

Time-scheduled operation with deterministic, chatter-free switching.

Scheduling is governed by a deterministic timetoindex map using the mid-points of the synthesis instants stored in `S.t`. Each interval \mathcal{I}_j activates model j until crossing the next midpoint; a dwell counter prevents spurious toggling. This choice mirrors the mission timeline and immediately induces an altitude-to-index map via $h(t)$. A dedicated segmented figure highlights which model is active over time: colored output segments,

black dots at synthesis instants, and light-grey vertical dashed lines at the switching boundaries. This visual device proved invaluable to interpret transition behavior and verify that switching is bumpless in all three controller variants.

Boundary compensation via K_{comp} and selective locking.

We introduced a simple yet effective compensation path: a scalar K_{comp} reweights the servo channel so that, at steady state, $r \approx K_{\text{comp}} y$. In *Nominal* operation, K_{comp} is estimated once and persisted per model (`S.ctrl(axis).Kcomp{j}`). In *Weighted* mode, the learned K_{comp} is applied at run time to reduce residual offsets at boundaries. We also added the option to (i) adapt K_{comp} online after switches using a lowpass ratio \bar{r}/\bar{y} with boost and bounds, and (ii) *lock* selected indices (e.g., enforce $K_{\text{comp}}(1) = 1$) to preserve design conventions in anchor regimes. Together, these features give a pragmatic knob to mitigate intermodel mismatch without altering the core LQG/LTR structure.

Predictive lookahead for sharper tracking.

The *Predictive* variant advances the reference by d samples inside the servo path. As shown in the comparative figures, this reduces lag and integral error around sharp trajectory features and disturbance entries. The expected trade-off, larger command activity, is visible in the control plots and quantified by $\sum(\Delta u_k)^2 T_s$ and J_{LQR} .

Robust wind disturbance rejection.

We injected representative wind disturbances and Gaussian measurement noise. In the three strategies, the controller remained robust up to the end of the run: no saturation, no runaway behavior, and bounded command. A small, bounded ripple appears in the transonic segment immediately after the first switch precisely where the plant varies fastest and where, by design, we *did not* apply gain interpolation. Far from being detrimental, this ripple is informative: it highlights where transition smoothing would yield the largest benefit while keeping the already adequate robustness margins.

A consistent set of metrics and reporting.

We formalized performance assessment with standard tracking indices (IAE/ISE/ITAE/ITSE), RMSE, effort terms $\sum u_k^2 T_s$, $\sum(\Delta u_k)^2 T_s$, peak command, and the discrete-time LQR cost $J_{\text{LQR}} = \sum_k (x_a^\top Q x_a + \Delta u^\top R \Delta u) T_s$. Metrics are reported *globally* and *per scheduled segment* using the same Q/R convention as in the controller. This made it possible to move from visual impressions to principled, quantitative comparisons: *Predictive* tends to minimize integral tracking indices; *Weighted* systematically reduces post-switch biases

and often improves RMSE/ISE; *Nominal* serves as a conservative baseline with the lowest actuation but larger residual errors. All conclusions were drawn on the pitch axis and carry over to roll and yaw *mutatis mutandis*, since synthesis and scheduling share the same structure and only the calibrated matrices differ.

Reproducibility and clarity.

The code base consolidates helper functions (normalization of D , robust P_c , state-estimation reconstruction from augmented states), plotting wrappers, and scheduling utilities. Inputs (reference, disturbances, $\mathbf{S.t}$) are validated and dimensioned at load time; all arrays are shaped to $1 \times N$ and trimmed/padded as needed. This careful engineering reduced integration errors and made the ablation between Nominal/Predictive/Weighted immediate.

In short, this dissertation delivered an operational, insight-oriented implementation of LQG/LTR for scheduled flight along a descent profile, complete with boundary compensation, predictive lookahead, disturbance injection, and rigorous metrics. The figures and tables consistently show robust behavior, meaningful trade-offs across strategies, and clear directions for refinement in the most demanding parts of the envelope (transonic/high q).

5.2 Future Works

The present framework opens several technically coherent and high-impact avenues:

(1) Smooth gain interpolation and bumpless transfer.

The observed boundary ripple in the transonic corridor strongly motivates interpolating scheduled gains— K , L , K_{servo} , P_c , and K_{comp} —either in time, altitude, Mach, or dynamic pressure. Classical choices include linear or spline interpolation with hysteresis/dwell; more refined options adopt LPV-style scheduling surfaces. A dedicated bumpless-transfer mechanism (e.g., command/observer bias scheduling) would further suppress boundary steps.

(2) Envelope-aware scheduling variables.

While time-based intervals are robust and simple, physically meaningful schedulers (Mach M , dynamic pressure q , altitude h) are better aligned with the true loci of plant variation. Combining $\mathbf{S.t}$ with $M(t)$ or $q(t)$ allows nonuniform partitions densified near transonic and high- q regimes. This can reduce the number of points while preserving fidelity where needed most.

(3) Structured robustness and multiobjective tuning.

The LTR step already recovers a loop shape with margin. Future work could (i) enforce structured robustness via H_∞/μ -synthesis on a small uncertainty template representative of aero-propulsive variation, or (ii) embed multi-objective tuning (e.g., Pareto optimization on IAE/ISE and $\sum(\Delta u)^2 T_s$) to formalize the trade between tracking and activity. Automated weight selection for Q/R over segments would standardize tuning across missions.

(4) Enhanced observers and noise shaping.

Switching-aware observers (scheduled Kalman gains, covariance propagation across boundaries) and noise-shaping (colored measurement noise, gust models) would increase realism. Nonlinear observers (EKF/UKF) could be explored to absorb mild nonlinearities near the boundary without abandoning the linear scheduled framework.

(5) Actuator constraints and anti-windup.

Although no saturation occurred here, future campaigns should explicitly include actuator limits and rate constraints, with an incremental anti-windup mechanism consistent with the augmented integrator. This will close the gap to hardware and prepare real-time deployments.

(6) Adaptive refinement of K_{comp} .

The online ratio-based update can be extended with (i) context-dependent filters (shorter time constants near boundaries), (ii) projection with soft bounds to avoid overshoots, and (iii) forgetting factors tied to flight condition. A small *a posteriori* bias observer, scheduled with the same intervals, would offer an alternative to scalar compensation.

(7) Higherfidelity plant and flexible modes.

Introducing aero tables varying with (M, q, h) , engine-gimbal dynamics, and low-order structural flexibility will stress the controller in ways that are typical of RLVs. The scheduled LQG/LTR core is expected to carry over with adjusted Q/R and possibly modest order augmentation.

(8) Hardwareintheloop and cyberphysical emulation.

A natural next step is to port the controller to a quadrotor-based cyber-physical emulator of descent dynamics (instrumentation, state estimation, and scheduled control

on-board), closing the loop between numerical validation and laboratory experiments. This will exercise timing, quantization, sensor faults, and task scheduling-dimensions that are hard to reproduce offline but crucial for eventual scalability.

(9) Software engineering and reproducibility.

Packaging `SIM_v02` as a self-contained toolbox (tests, examples, data loaders, export of LaTeXready figures/tables) would make the workflow immediately reusable. A small CLI or `live script` could let users select axes, strategies, and scheduling options, then produce plots and metrics in one shot. Versioned configuration files would ensure consistent comparisons across runs.

(10) From pitch exemplar to fullaxis integration.

All analyses were presented on the pitch axis as a representative case; rolling and yawing axes share the same synthesis and scheduling. A full-axis integration-possibly with cross-axis couplings-would consolidate the methodology and reveal multi-axis interactions at boundaries.

Overall, the present dissertation demonstrates that a disciplined LQG/LTR design, paired with transparent scheduling and simple boundary compensation, can deliver robust tracking throughout aerodynamic descent, even under wind disturbances. The framework is deliberately modest-linear, scheduled, and diagnostic-rich-so that refinements (gain interpolation, structured robustness, hardware emulation) can be layered without redesigning the core. This is, we believe, a solid foundation for progressing from reliable simulations to confident experiments and, ultimately, to scalable solutions for reusable launch systems.

Appendix

Since the angle ϕ is very close to 180 deg, and θ and ψ are close to 0 deg, Eq. (3) becomes the following:

$$\mathbf{R}_B^{\text{DCA}} \cong \begin{bmatrix} 1 & \psi & -\theta \\ \psi & -1 & \phi - \pi \\ -\theta & \pi - \phi & -1 \end{bmatrix} \quad (\text{A1})$$

By plugging Eq. (5), Eq. (6), and Eq. (A1) into Eq. (1), and by expressing the equations componentwise we obtain

$$\begin{aligned} \dot{A} &= v_A \\ \dot{C} &= v_C \\ \dot{D} &= v_D \\ \dot{v}_A &= \mathbf{a}_A^{\text{grav}} + \frac{1}{2}\rho \frac{(v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2)S}{m} \\ &\quad \times C_A(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi) + \mathbf{a}_A^{\text{fict}} \\ \dot{v}_C &= \mathbf{a}_C^{\text{grav}} + \frac{1}{2}\rho \frac{(v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2)S}{m} \\ &\quad \times C_C(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi) + \mathbf{a}_C^{\text{fict}} \\ \dot{v}_D &= \mathbf{a}_D^{\text{grav}} + \frac{1}{2}\rho \frac{(v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2)S}{m} \\ &\quad \times C_D(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi) + \mathbf{a}_D^{\text{fict}} \\ \dot{\phi} &= \omega_x \\ \dot{\theta} &= -\omega_y \\ \dot{\psi} &= -\omega_z \\ \dot{\omega}_x &= \frac{1}{2}\rho \frac{(v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2)SL}{I_{xx}} \\ &\quad \times C_l(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi) + I_{\omega,x} \\ \dot{\omega}_y &= \frac{1}{2}\rho \frac{(v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2)SL}{I_{yy}} \\ &\quad \times C_m(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi) + I_{\omega,y} \\ \dot{\omega}_z &= \frac{1}{2}\rho \frac{(v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2)SL}{I_{zz}} \\ &\quad \times C_n(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi) + I_{\omega,z} \end{aligned} \quad (\text{A2})$$

with

$$\begin{aligned} C_A &= C_x + C_y\psi - C_z\theta \\ C_C &= C_x\psi - C_y + C_z(\phi - \pi) \\ C_D &= -C_x\theta + C_y(\pi - \phi) - C_z \end{aligned} \quad (\text{A3})$$

coming from the application of Eq. (A1) to the aerodynamic coefficients with respect to the body reference frame C_x , C_y , and C_z . The terms $I_{\omega,x}$, $I_{\omega,y}$, and $I_{\omega,z}$ are the components of the vector $\boldsymbol{\omega} \times \mathbf{I} \cdot \boldsymbol{\omega}$, neglected from now on, as always smaller than 1 [N · m]. Similarly, we neglect the fictitious terms $\mathbf{a}_A^{\text{fict}}$, $\mathbf{a}_C^{\text{fict}}$, and $\mathbf{a}_D^{\text{fict}}$ as they are always much smaller than the other contributions to the accelerations. As previously mentioned, the wind is assumed to act only in the local horizontal plane; that is, $\mathbf{v}_w = [0, v_{w,C}, v_{w,D}]^T$.

Equation (A2) can be linearized around the reference solution. To obtain the linearized forms needed for the control synthesis we have to compute how each of these differential equations changes when the corresponding independent variables change. Note that formally we are defining perturbations with respect to the reference solutions; therefore a Δ symbol is here adopted. However, the symbol is omitted in the main text for avoiding heavy notation, with the clear intention to refer to the variations of the variables with respect to their reference value rather than to the variables themselves.

A.1. Roll

For the roll angle we have

$$\begin{aligned} \Delta \dot{\phi} &= \Delta \omega_x \\ \Delta \dot{\omega}_x &= \rho \frac{(SL)}{I_{xx}} C_l [v_A \Delta v_A + v_C \Delta v_C + v_D \Delta v_D] + \dots \\ &\quad + \frac{1}{2}\rho \frac{(v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2)SL}{I_{xx}} \\ &\quad \times [C_{l,\phi} \Delta \phi + C_{l,\theta} \Delta \theta + C_{l,\psi} \Delta \psi + C_{l,\delta_\phi} \Delta \delta_\phi + C_{l,\delta_\theta} \Delta \delta_\theta \\ &\quad + C_{l,\delta_\psi} \Delta \delta_\psi] \end{aligned} \quad (\text{A4})$$

Since the nominal roll coefficient C_l is zero, the first term on the right-hand side of the $\Delta \dot{\omega}_x$ equation containing the contributions of Δv_A , Δv_C , and Δv_D disappears. By numerical inspection of the derivatives $C_{l,\phi}$, $C_{l,\theta}$, $C_{l,\psi}$, C_{l,δ_ϕ} , C_{l,δ_θ} , and C_{l,δ_ψ} the following condition holds:

$$C_{l,\delta_\phi} \gg C_{l,\phi}, C_{l,\theta}, C_{l,\psi}, C_{l,\delta_\theta}, C_{l,\delta_\psi} \quad (\text{A5})$$

which means that the roll fin deflection dominates the motion, and therefore the other terms, representing the coupling between axes, can be neglected, making the motion of roll similar to a double integrator. However, while the terms $C_{l,\theta}$, $C_{l,\psi}$, C_{l,δ_θ} , and C_{l,δ_ψ} are assumed to be 0, the term $C_{l,\phi}$ is retained, because it can easily be incorporated in the corresponding state-space representation. By defining $a_0 = \frac{q_{\text{dyn}} S_{\text{ref}} L_{\text{ref}}}{I_{xx}}$ we retrieve the state-space expressions of Eq. (8).

A.2. Pitch

Let us consider again Eq. (A2), and specifically the differential equations for D , v_D , θ , and ω_y .

$$\begin{aligned} \dot{D} &= v_D \\ \dot{v}_D &= \mathbf{a}_D^{\text{grav}} + \frac{1}{2}\rho \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} \\ &\quad \times C_D(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi) \\ \dot{\theta} &= -\omega_y \\ \dot{\omega}_y &= \frac{1}{2}\rho \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]SL}{I_{yy}} \\ &\quad \times C_m(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi) \end{aligned} \quad (\text{A6})$$

The first and the third equations are linear and therefore require no further work, as the corresponding variations will be clear.

$$\begin{aligned} \Delta \dot{D} &= \Delta v_D \\ \Delta \dot{\theta} &= -\Delta \omega_y \end{aligned} \quad (\text{A7})$$

Let us focus on the \dot{v}_D equation: by differentiating with respect to D , C , v_D , v_C , ϕ , θ , ψ , δ_ϕ , δ_θ , δ_ψ , $v_{w,C}$, and $v_{w,D}$, we get

$$\begin{aligned} \Delta \dot{v}_D &= \frac{\partial \dot{v}_D}{\partial A} \Delta A + \frac{\partial \dot{v}_D}{\partial C} \Delta C + \frac{\partial \dot{v}_D}{\partial D} \Delta D + \frac{\partial \dot{v}_D}{\partial v_A} \Delta v_A + \frac{\partial \dot{v}_D}{\partial v_C} \Delta v_C \\ &\quad + \frac{\partial \dot{v}_D}{\partial v_D} \Delta v_D + \frac{\partial \dot{v}_D}{\partial \phi} \Delta \phi + \frac{\partial \dot{v}_D}{\partial \theta} \Delta \theta + \frac{\partial \dot{v}_D}{\partial \psi} \Delta \psi + \frac{\partial \dot{v}_D}{\partial \delta_\phi} \Delta \delta_\phi \\ &\quad + \frac{\partial \dot{v}_D}{\partial \delta_\theta} \Delta \delta_\theta + \frac{\partial \dot{v}_D}{\partial \delta_\psi} \Delta \delta_\psi + \frac{\partial \dot{v}_D}{\partial v_{w,D}} \Delta v_{w,D} + \frac{\partial \dot{v}_D}{\partial v_{w,C}} \Delta v_{w,C} \end{aligned} \quad (\text{A8})$$

Since the motion of the rocket is not powered, the altitude and the altitude rate are not considered explicitly in the state-space representation. This choice is also justified by the use of altitude as scheduling parameter. Let us compute the other derivatives. The dependency on position components is very small, and only due to minor changes in

the local altitude caused by a change in downrange/crossrange as they are defined with respect to the landing site. By taking into account Eq. (2) we get

$$\begin{aligned} \frac{\partial \dot{v}_D}{\partial D} &= \frac{\partial \mathbf{a}_D^{\text{grav}}}{\partial D} + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{D,M} \frac{\partial M}{\partial D} \\ &= \mu_{\oplus} \frac{3(D + D_L)^2}{[(A + A_L)^2 + (C + C_L)^2 + (D + D_L)^2]^{\frac{3}{2}}} \\ &\quad - \frac{1}{[(A + A_L)^2 + (C + C_L)^2 + (D + D_L)^2]^{\frac{3}{2}}} \\ &\quad + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{D,M} \frac{\partial M}{\partial D} \\ \frac{\partial \dot{v}_C}{\partial C} &= \frac{\partial \mathbf{a}_C^{\text{grav}}}{\partial C} + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{C,M} \frac{\partial M}{\partial C} \\ &= \mu_{\oplus} \frac{3(C + C_L)^2}{[(A + A_L)^2 + (C + C_L)^2 + (D + D_L)^2]^{\frac{3}{2}}} \\ &\quad - \frac{1}{[(A + A_L)^2 + (C + C_L)^2 + (D + D_L)^2]^{\frac{3}{2}}} \\ &\quad + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{C,M} \frac{\partial M}{\partial C} \\ \frac{\partial \dot{v}_D}{\partial C} &= \frac{\partial \mathbf{a}_D^{\text{grav}}}{\partial C} + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{D,M} \frac{\partial M}{\partial C} \\ &= \mu_{\oplus} \frac{3(D + D_L)(C + C_L)}{[(A + A_L)^2 + (C + C_L)^2 + (D + D_L)^2]^{\frac{3}{2}}} \\ &\quad + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{D,M} \frac{\partial M}{\partial C} \\ \frac{\partial \dot{v}_C}{\partial D} &= \frac{\partial \mathbf{a}_D^{\text{grav}}}{\partial D} + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{C,M} \frac{\partial M}{\partial D} \\ &= \mu_{\oplus} \frac{3(D + D_L)(C + C_L)}{[(A + A_L)^2 + (C + C_L)^2 + (D + D_L)^2]^{\frac{3}{2}}} \\ &\quad + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{C,M} \frac{\partial M}{\partial D} \quad (\text{A9}) \end{aligned}$$

For the dependency on the velocity components v_D and v_C , by looking at Eq. (A6) we can write

$$\begin{aligned} \frac{\partial \dot{v}_D}{\partial v_D} &= \rho \frac{S}{m} (v_D + v_{w,D}) C_D \\ &\quad + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{D,M} \frac{\partial M}{\partial v} \frac{\partial v}{\partial v_D} \\ \frac{\partial \dot{v}_C}{\partial v_C} &= \rho \frac{S}{m} (v_C + v_{w,C}) C_C \\ &\quad + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{C,M} \frac{\partial M}{\partial v} \frac{\partial v}{\partial v_C} \quad (\text{A10}) \end{aligned}$$

where $C_{D,M}$ and $C_{C,M}$ are the partial derivatives of the coefficients C_D and C_C with respect to the Mach number M , obtained numerically from the aerodynamic database, whereas the terms $\frac{\partial M}{\partial v}$, $\frac{\partial v}{\partial v_D}$, and $\frac{\partial v}{\partial v_C}$ are computed as follows:

$$\begin{aligned} \frac{\partial M}{\partial v} &= \frac{1}{\sqrt{\gamma_{\text{air}} R_{\text{air}} T}} \\ \frac{\partial v}{\partial v_D} &= \frac{\partial [v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]^{\frac{1}{2}}}{\partial v_D} = \frac{v_D + v_{w,D}}{v} \\ \frac{\partial v}{\partial v_C} &= \frac{\partial [v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]^{\frac{1}{2}}}{\partial v_C} = \frac{v_C + v_{w,C}}{v} \quad (\text{A11}) \end{aligned}$$

where γ_{air} and R_{air} the specific heat ratio of the air, and the specific air constant, assumed equal to 1.4 and $287.058 \text{ J} \cdot \text{kg}^{-1} \cdot \text{K}^{-1}$, respectively, while T is the atmospheric temperature in K, and the implicit definition $v = [v_A^2 + (v_D + v_{w,D})^2 + (v_C + v_{w,C})^2]^{\frac{1}{2}}$. Note that we derived also some terms associated with \dot{v}_C , which will be useful later. Let us consider the dependency on the attitude angles ϕ , θ , and ψ , and the virtual fin deflections δ_ϕ , δ_θ , and δ_ψ . In this case the expressions are immediate, as they are driven by the numerical derivatives extracted from the aerodynamic database, combined with Eq. (A3):

$$\begin{aligned} \frac{\partial \dot{v}_D}{\partial \phi} &= \frac{1}{2\rho} \frac{v^2 S}{m} [-C_{x,\phi} \theta + C_{y,\phi} (\pi - \phi) - C_y - C_{z,\phi}] \\ \frac{\partial \dot{v}_D}{\partial \theta} &= \frac{1}{2\rho} \frac{v^2 S}{m} [-C_{x,\theta} \theta - C_x + C_{y,\theta} (\pi - \phi) - C_{z,\theta}] \\ \frac{\partial \dot{v}_D}{\partial \psi} &= \frac{1}{2\rho} \frac{v^2 S}{m} [-C_{x,\psi} \theta + C_{y,\psi} (\pi - \phi) - C_{z,\psi}] \\ \frac{\partial \dot{v}_D}{\partial \delta_\phi} &= \frac{1}{2\rho} \frac{v^2 S}{m} [-C_{x,\delta_\phi} \theta + C_{y,\delta_\phi} (\pi - \phi) - C_{z,\delta_\phi}] \\ \frac{\partial \dot{v}_D}{\partial \delta_\theta} &= \frac{1}{2\rho} \frac{v^2 S}{m} [-C_{x,\delta_\theta} \theta + C_{y,\delta_\theta} (\pi - \phi) - C_{z,\delta_\theta}] \\ \frac{\partial \dot{v}_D}{\partial \delta_\psi} &= \frac{1}{2\rho} \frac{v^2 S}{m} [-C_{x,\delta_\psi} \theta + C_{y,\delta_\psi} (\pi - \phi) - C_{z,\delta_\psi}] \quad (\text{A12}) \end{aligned}$$

The remaining derivatives to be analyzed are the ones with respect to the wind components $v_{w,D}$ and $v_{w,C}$. Given the structure of the equations, it is immediate to verify that they are simply given by

$$\begin{aligned} \frac{\partial \dot{v}_D}{\partial v_{w,D}} &= \frac{\partial \dot{v}_D}{\partial v_D} \\ \frac{\partial \dot{v}_D}{\partial v_{w,C}} &= \frac{\partial \dot{v}_D}{\partial v_C} \quad (\text{A13}) \end{aligned}$$

For what regards the angular rate ω_y , we can compute the first-order variation with respect to the reference motion as

$$\begin{aligned} \Delta \dot{\omega}_y &= \frac{\partial \dot{\omega}_y}{\partial A} \Delta A + \frac{\partial \dot{\omega}_y}{\partial C} \Delta C + \frac{\partial \dot{\omega}_y}{\partial D} \Delta D + \frac{\partial \dot{\omega}_y}{\partial v_A} \Delta v_A + \frac{\partial \dot{\omega}_y}{\partial v_C} \Delta v_C \\ &\quad + \frac{\partial \dot{\omega}_y}{\partial v_D} \Delta v_D + \frac{\partial \dot{\omega}_y}{\partial \phi} \Delta \phi + \frac{\partial \dot{\omega}_y}{\partial \theta} \Delta \theta + \frac{\partial \dot{\omega}_y}{\partial \psi} \Delta \psi + \frac{\partial \dot{\omega}_y}{\partial \delta_\phi} \Delta \delta_\phi \\ &\quad + \frac{\partial \dot{\omega}_y}{\partial \delta_\theta} \Delta \delta_\theta + \frac{\partial \dot{\omega}_y}{\partial \delta_\psi} \Delta \delta_\psi + \frac{\partial \dot{\omega}_y}{\partial v_{w,D}} \Delta v_{w,D} + \frac{\partial \dot{\omega}_y}{\partial v_{w,C}} \Delta v_{w,C} \quad (\text{A14}) \end{aligned}$$

The dependencies of the angular rate ω_y on position are reflected in the following expressions:

$$\begin{aligned} \frac{\partial \dot{\omega}_y}{\partial D} &= \frac{1}{2} \frac{\partial \rho}{\partial D} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_m \\ &\quad + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{m,M} \frac{\partial M}{\partial D} \\ \frac{\partial \dot{\omega}_y}{\partial C} &= \frac{1}{2} \frac{\partial \rho}{\partial C} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_m \\ &\quad + \frac{1}{2\rho} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2]S}{m} C_{m,M} \frac{\partial M}{\partial C} \quad (\text{A15}) \end{aligned}$$

with the terms $\frac{\partial \rho}{\partial D}$ and $\frac{\partial M}{\partial D}$ computed numerically, whereas the dependency on the velocity components is captured by the following expressions:

$$\begin{aligned}\frac{\partial \dot{\omega}_y}{\partial v_D} &= \rho \frac{SL}{I_{yy}} (v_D + v_{w,D}) C_m \\ &+ \frac{1}{2} \rho \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2] SL}{I_{yy}} C_{m,M} \frac{\partial M}{\partial v} \frac{\partial v}{\partial v_D} \\ \frac{\partial \dot{\omega}_y}{\partial v_C} &= \rho \frac{SL}{I_{yy}} (v_C + v_{w,C}) C_m \\ &+ \frac{1}{2} \rho \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2] SL}{I_{yy}} C_{m,M} \frac{\partial M}{\partial v} \frac{\partial v}{\partial v_C}\end{aligned}\quad (\text{A16})$$

Similarly, for the attitude and the fin deflection dependencies, we have

$$\begin{aligned}\frac{\partial \dot{\omega}_y}{\partial \phi} &= \frac{1}{2} \rho v^2 \frac{SL}{I_{yy}} C_{m,\phi}, & \frac{\partial \dot{\omega}_y}{\partial \theta} &= \frac{1}{2} \rho v^2 \frac{SL}{I_{yy}} C_{m,\theta}, \\ \frac{\partial \dot{\omega}_y}{\partial \psi} &= \frac{1}{2} \rho v^2 \frac{SL}{I_{yy}} C_{m,\psi} \\ \frac{\partial \dot{\omega}_y}{\partial \delta_\phi} &= \frac{1}{2} \rho v^2 \frac{SL}{I_{yy}} C_{m,\delta_\phi}, & \frac{\partial \dot{\omega}_y}{\partial \delta_\theta} &= \frac{1}{2} \rho v^2 \frac{SL}{I_{yy}} C_{m,\delta_\theta}, \\ \frac{\partial \dot{\omega}_y}{\partial \delta_\psi} &= \frac{1}{2} \rho v^2 \frac{SL}{I_{yy}} C_{m,\delta_\psi}\end{aligned}\quad (\text{A17})$$

Finally, the dependency on wind is captured by the following expressions:

$$\frac{\partial \dot{\omega}_y}{\partial v_{w,D}} = \frac{\partial \dot{\omega}_y}{\partial v_D}, \quad \frac{\partial \dot{\omega}_y}{\partial v_{w,C}} = \frac{\partial \dot{\omega}_y}{\partial v_C}\quad (\text{A18})$$

By inspection of all the terms in Eqs. (A8) and (A14) it is possible to see that the following conditions hold.

$$\frac{\partial \dot{\omega}_y}{\partial v_D} \gg \frac{\partial \dot{\omega}_y}{\partial v_C}\quad (\text{A19})$$

which implies that

$$\frac{\partial \dot{\omega}_y}{\partial v_{w,D}} \gg \frac{\partial \dot{\omega}_y}{\partial v_{w,C}}\quad (\text{A20})$$

Moreover, by looking at Eq. (A17), the following conditions are satisfied:

$$\begin{aligned}\frac{\partial \dot{v}_D}{\partial v_D} &\gg \frac{\partial \dot{v}_D}{\partial C}, \frac{\partial \dot{v}_D}{\partial D}, \frac{\partial \dot{v}_D}{\partial v_C} \\ \frac{\partial \dot{v}_D}{\partial \theta} &\gg \frac{\partial \dot{v}_D}{\partial \phi}, \frac{\partial \dot{v}_D}{\partial \psi} \\ C_{m,v_D} &\gg C_{m,v_C} \\ C_{m,\theta} &\gg C_{m,\phi}, C_{m,\psi} \\ C_{m,\delta_\theta} &\gg C_{m,\delta_\phi}, C_{m,\delta_\psi}\end{aligned}\quad (\text{A21})$$

and therefore the cross-coupling terms can effectively be neglected. However, also in this case, we keep the term $\frac{\partial \dot{v}_D}{\partial D}$ as it does not add any complexity in the modeling of the corresponding state-space representation. By defining $a_1 = \frac{q_{dyn} S}{m}$ and $a_2 = \frac{q_{dyn} SL}{I_{yy}}$ and by introducing

$$\begin{aligned}a_{v_D,D} &= \frac{\partial a_D^{grav}}{\partial D} \\ a_{v_D,v_D} &= \rho \frac{S}{m} (v_D + v_{w,D}) C_D + a_1 C_{D,M} \frac{\partial M}{\partial v} \frac{\partial v}{\partial v_D} \\ a_{v_D,\theta} &= a_1 \frac{1}{2} \rho \frac{v^2 S}{m} [-C_{x,\theta} \theta - C_x + C_{y,\theta} (\pi - \phi) - C_{z,\theta}] \\ b_{v_D,\delta_\theta} &= a_1 \frac{1}{2} \rho \frac{v^2 S}{m} [-C_{x,\delta_\theta} \theta + C_{y,\delta_\theta} (\pi - \phi) - C_{z,\delta_\theta}] \\ b_{v_D,v_{w,D}} &= a_{v_D,v_D} \\ a_{\omega_y,D} &= \frac{1}{2} \frac{\partial \rho}{\partial D} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2] S}{m} C_m \\ &+ a_1 C_{m,M} \frac{\partial M}{\partial D} \\ a_{\omega_y,v_D} &= \rho \frac{SL}{I_{yy}} (v_D + v_{w,D}) C_m + a_2 C_{m,M} \frac{\partial M}{\partial v} \frac{\partial v}{\partial v_D} \\ b_{\omega_y,\delta_\theta} &= a_2 C_{m,\delta_\theta} \\ b_{\omega_y,v_{w,D}} &= a_{\omega_y,v_D}\end{aligned}\quad (\text{A22})$$

we obtain the system described by Eq. (10).

A.3. Yaw

Finally, for completeness, let us provide the results for yaw. The corresponding equations to be analyzed are

$$\begin{aligned}\dot{C} &= v_C \\ \dot{v}_C &= a_C^{grav} + \frac{1}{2} \rho \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2] S}{m} \\ &\quad \times C_C(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi) \\ \dot{\psi} &= -\omega_z \\ \dot{\omega}_z &= \frac{1}{2} \rho \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2] SL}{I_{zz}} \\ &\quad \times C_n(\phi, \theta, \psi, M, \delta_\phi, \delta_\theta, \delta_\psi)\end{aligned}\quad (\text{A23})$$

Also in this case the first and the third equations are linear.

$$\begin{aligned}\Delta \dot{C} &= \Delta v_C \\ \Delta \dot{\psi} &= -\Delta \omega_z\end{aligned}\quad (\text{A24})$$

The differentiation of \dot{v}_C with respect to D , C , v_D , v_C , ϕ , θ , ψ , δ_ϕ , δ_θ , δ_ψ , $v_{w,C}$, and $v_{w,D}$ gives

$$\begin{aligned}\Delta \dot{v}_C &= \frac{\partial \dot{v}_C}{\partial A} \Delta A + \frac{\partial \dot{v}_C}{\partial C} \Delta C + \frac{\partial \dot{v}_C}{\partial D} \Delta D + \frac{\partial \dot{v}_C}{\partial v_A} \Delta v_A + \frac{\partial \dot{v}_C}{\partial v_C} \Delta v_C \\ &\quad + \frac{\partial \dot{v}_C}{\partial v_D} \Delta v_D + \frac{\partial \dot{v}_C}{\partial \phi} \Delta \phi + \frac{\partial \dot{v}_C}{\partial \theta} \Delta \theta + \frac{\partial \dot{v}_C}{\partial \psi} \Delta \psi + \frac{\partial \dot{v}_C}{\partial \delta_\phi} \Delta \delta_\phi \\ &\quad + \frac{\partial \dot{v}_C}{\partial \delta_\theta} \Delta \delta_\theta + \frac{\partial \dot{v}_C}{\partial \delta_\psi} \Delta \delta_\psi + \frac{\partial \dot{v}_C}{\partial v_{w,D}} \Delta v_{w,D} + \frac{\partial \dot{v}_C}{\partial v_{w,C}} \Delta v_{w,C}\end{aligned}\quad (\text{A25})$$

The derivatives of interest with respect to the velocity components v_D and v_C can be retrieved from Eqs. (A9–A11). For the attitude and the fin deflections dependencies [and again, taking Eq. (A3) into account], we obtain

$$\begin{aligned}
\frac{\partial \dot{v}_C}{\partial \phi} &= \frac{1}{2} \rho \frac{v^2 S}{m} [C_{x,\phi} \psi - C_{y,\phi} + C_{z,\phi}(\phi - \pi) + C_z] \\
\frac{\partial \dot{v}_C}{\partial \theta} &= \frac{1}{2} \rho \frac{v^2 S}{m} [C_{x,\theta} \psi - C_{y,\theta} + C_{z,\theta}(\phi - \pi)] \\
\frac{\partial \dot{v}_C}{\partial \psi} &= \frac{1}{2} \rho \frac{v^2 S}{m} [C_{x,\psi} \psi + C_x - C_{y,\psi} + C_{z,\psi}(\phi - \pi)] \\
\frac{\partial \dot{v}_C}{\partial \delta_\phi} &= \frac{1}{2} \rho \frac{v^2 S}{m} [C_{x,\delta_\phi} \psi - C_{y,\delta_\phi} + C_{z,\delta_\phi}(\phi - \pi)] \\
\frac{\partial \dot{v}_C}{\partial \delta_\theta} &= \frac{1}{2} \rho \frac{v^2 S}{m} [C_{x,\delta_\theta} \theta - C_{y,\delta_\theta} + C_{z,\delta_\theta}(\phi - \pi)] \\
\frac{\partial \dot{v}_C}{\partial \delta_\psi} &= \frac{1}{2} \rho \frac{v^2 S}{m} [C_{x,\delta_\psi} \theta - C_{y,\delta_\psi} + C_{z,\delta_\psi}(\phi - \pi)] \quad (A26)
\end{aligned}$$

Finally, in parallel with what shown for Eq. (A13), we get

$$\begin{aligned}
\frac{\partial \dot{v}_C}{\partial v_{w,D}} &= \frac{\partial \dot{v}_C}{\partial v_D} \\
\frac{\partial \dot{v}_C}{\partial v_{w,C}} &= \frac{\partial \dot{v}_C}{\partial v_C} \quad (A27)
\end{aligned}$$

The procedure is exactly the same for the angular rate component $\dot{\omega}_z$, for which we can find

$$\begin{aligned}
\frac{\partial \dot{\omega}_z}{\partial v_D} &= \rho \frac{SL}{I_{zz}} (v_D + v_{w,D}) C_n \\
&\quad + \frac{1}{2} \rho \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2] SL}{I_{zz}} C_{n,M} \frac{\partial M}{\partial v} v_{v_D} \\
\frac{\partial \dot{\omega}_z}{\partial v_C} &= \rho \frac{SL}{I_{zz}} (v_C + v_{w,C}) C_n \\
&\quad + \frac{1}{2} \rho \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2] SL}{I_{zz}} C_{n,M} \frac{\partial M}{\partial v} v_{v_C} \quad (A28)
\end{aligned}$$

for the velocity components,

$$\begin{aligned}
\frac{\partial \dot{\omega}_z}{\partial \phi} &= \frac{1}{2} \rho \frac{SL}{I_{zz}} C_{n,\phi}, & \frac{\partial \dot{\omega}_z}{\partial \theta} &= \frac{1}{2} \rho \frac{SL}{I_{zz}} C_{n,\theta} \\
\frac{\partial \dot{\omega}_z}{\partial \psi} &= \frac{1}{2} \rho \frac{SL}{I_{zz}} C_{n,\psi}, \\
\frac{\partial \dot{\omega}_z}{\partial \delta_\phi} &= \frac{1}{2} \rho \frac{SL}{I_{zz}} C_{n,\delta_\phi}, \\
\frac{\partial \dot{\omega}_z}{\partial \delta_\theta} &= \frac{1}{2} \rho \frac{SL}{I_{zz}} C_{n,\delta_\theta}, \\
\frac{\partial \dot{\omega}_z}{\partial \delta_\psi} &= \frac{1}{2} \rho \frac{SL}{I_{zz}} C_{n,\delta_\psi} \quad (A29)
\end{aligned}$$

for the attitude and the fin deflections terms, and

$$\frac{\partial \dot{\omega}_z}{\partial v_{w,D}} = \frac{\partial \dot{\omega}_z}{\partial v_D}, \quad \frac{\partial \dot{\omega}_z}{\partial v_{w,C}} = \frac{\partial \dot{\omega}_z}{\partial v_C} \quad (A30)$$

for the wind. The observations made for pitch hold here too. The yaw-related terms dominate the behavior, or in other words

$$\begin{aligned}
\frac{\partial \dot{v}_C}{\partial v_C} &\gg \frac{\partial \dot{v}_C}{\partial D}, \frac{\partial \dot{v}_C}{\partial C}, \frac{\partial \dot{v}_C}{\partial v_D} \\
\frac{\partial \dot{v}_C}{\partial \psi} &\gg \frac{\partial \dot{v}_C}{\partial \phi}, \frac{\partial \dot{v}_C}{\partial \theta} \\
C_{n,v_C} &\gg C_{n,v_D} \\
C_{n,\psi} &\gg C_{n,\phi}, C_{n,\theta} \\
C_{n,\delta_\psi} &\gg C_{n,\delta_\phi}, C_{n,\delta_\theta} \quad (A31)
\end{aligned}$$

By applying the definitions

$$\begin{aligned}
a_{v_C,C} &= \frac{\partial a_C^{\text{grav}}}{\partial C} \\
a_{v_C,v_C} &= \rho \frac{S}{m} (v_C + v_{w,C}) C_C + a_1 C_{C,M} \frac{\partial M}{\partial v} \frac{\partial v}{\partial v_C} \\
a_{v_C,\psi} &= a_1 \frac{1}{2} \rho \frac{v^2 S}{m} [-C_{x,\psi} \theta + C_{y,\psi}(\pi - \phi) - C_{z,\psi}] \\
b_{v_C,\delta_\psi} &= a_1 \frac{1}{2} \rho \frac{v^2 S}{m} [-C_{x,\delta_\psi} \theta + C_{y,\delta_\psi}(\pi - \phi) - C_{z,\delta_\psi}] \\
b_{v_C,v_{w,C}} &= a_{v_C,v_C} \\
a_{\omega_z,C} &= \frac{1}{2} \frac{\partial \rho}{\partial D} \frac{[v_A^2 + (v_C + v_{w,C})^2 + (v_D + v_{w,D})^2] S}{m} C_n \\
&\quad + a_3 C_{n,M} \frac{\partial M}{\partial C} \\
a_{\omega_z,v_C} &= \rho \frac{SL}{I_{zz}} (v_C + v_{w,C}) C_n + a_2 C_{n,M} \frac{\partial M}{\partial v} \frac{\partial v}{\partial v_C} \\
b_{\omega_z,\delta_\psi} &= a_2 C_{n,\delta_\psi} \\
b_{\omega_z,v_{w,C}} &= a_{\omega_z,v_C} \quad (A32)
\end{aligned}$$

we can finally build a corresponding state-space representation for yaw.

$$\mathbf{x} = [C \quad v_C \quad \psi \quad \omega_z]^T \quad (A33)$$

with the control $u = \delta_\psi$, and the following state-space realization:

$$\begin{aligned}
\begin{bmatrix} \dot{C} \\ \dot{v}_C \\ \dot{\psi} \\ \dot{\omega}_z \end{bmatrix} &= \begin{bmatrix} 0 & 1 & 0 & 0 \\ a_{v_C,C} & a_{v_C,v_C} & a_{v_C,\psi} & 0 \\ 0 & 0 & 0 & -1 \\ a_{\omega_z,C} & a_{\omega_z,v_C} & a_{\omega_z,\psi} & 0 \end{bmatrix} \begin{bmatrix} C \\ v_C \\ \psi \\ \omega_z \end{bmatrix} \\
&\quad + \begin{bmatrix} 0 & 0 \\ b_{v_C,\delta_\psi} & b_{v_C,v_{w,C}} \\ 0 & 0 \\ b_{\omega_z,\delta_\psi} & b_{\omega_z,v_{w,C}} \end{bmatrix} \begin{bmatrix} \delta_\psi \\ v_{w,C} \end{bmatrix} \quad (A34)
\end{aligned}$$

Appendix B: Reference Frames, Aerodynamic Angles, and Actuators Description

B.1. Reference Frames

1) *Vehicle reference frame*: The vehicle reference frame is a fixed system of axes centered in the gimbal point of the thrust vector control (TVC) of the rocket. The X_V axis points toward the nose of the rocket, whereas the Z_V axis is aligned with the plane of symmetry of leg 1. Y_V is given by $Z_V \times X_V$.

2) *Body reference frame*: The body reference frame of the rocket has its origin in the center of mass of the rocket itself, and the axes parallel with the vehicle reference frame. Since there is no mass consumption during the aerodynamic descent, the transformation between the body and the vehicle reference frames is given by a fixed translation along the X_B axis of about 5 m. The body reference frame is shown in Fig. 1a.

3) *Air-path reference frame*: It is defined according to the ISO norm 1151, Part 1, Sec. 1.1, and visible in Fig. B1.

B.2. Aerodynamic Angles

The angle of attack α and the sideslip angle are defined accordingly to the corresponding ISO norm (ISO 1151, Part 1, Sec 1.2), and are defined as follows:

Annex C: Smooth blending weights for the wind profile

We blend the staged speed laws with C^1 weights built from the canonical *smoothstep* window. For an altitude z and bounds $z_1 < z_2$:

$$s(z; z_1, z_2) = \frac{z - z_1}{\max(z_2 - z_1, \varepsilon)}, \quad \hat{s} = \text{clamp}_{[0,1]}(s), \quad \text{smoothstep}(z; z_1, z_2) = \hat{s}^2(3 - 2\hat{s}).$$

Thus, with the altitude bands (z_{12}^1, z_{12}^2) , (z_{23}^1, z_{23}^2) , (z_{34}^1, z_{34}^2) :

$$w_{12} = \text{smoothstep}(z; z_{12}^1, z_{12}^2), \quad w_{23} = \text{smoothstep}(z; z_{23}^1, z_{23}^2),$$

$$w_{34} = \text{smoothstep}(z; z_{34}^1, z_{34}^2).$$

Given the stage targets $U_{\log}(z)$, $U_{\text{pow}}(z)$, $U_{\text{tran}}(z)$, $U_{\text{jet}}(z)$, the smoothly blended speed is

$$U_{12} = (1 - w_{12})U_{\log} + w_{12}U_{\text{pow}}, \quad U_{23} = (1 - w_{23})U_{12} + w_{23}U_{\text{tran}},$$

$$U(z) = (1 - w_{34})U_{23} + w_{34}U_{\text{jet}}.$$

Direction convention (“TO”).

Let the base direction be ψ_{base} (`base_dir_deg`) and a boolean `dir_is_from`. The mean direction *to* is

$$\psi_w^{\text{TO}}(z) = \psi_{\text{base}} + 180 \mathbb{I}_{\text{FROM}} + \Delta\psi_{\text{max}} \left[1 - \exp\left(-\left(\frac{z}{z_v}\right)^p\right) \right] + \varepsilon \sin\left(\frac{z}{\ell}\right),$$

with typical $(\Delta\psi_{\text{max}}, z_v, p, \varepsilon, \ell) = (25^\circ, 1000 \text{ m}, 1.2, 5^\circ, 400 \text{ m})$. Components in NED then read $\bar{\mathbf{w}}_n = [U(z) \cos \psi_w^{\text{TO}}, U(z) \sin \psi_w^{\text{TO}}, 0]^\top$.

References

- ANDERSON, B. D. O.; MOORE, J. B. *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall, 1979. Foundational treatment of the discrete Kalman filter and Riccati equations. Cited 2 times on the pages 19 and 36.
- ÅSTRÖM, K. J.; WITTENMARK, B. *Computer-Controlled Systems: Theory and Design*. 3rd. ed. Upper Saddle River, NJ, USA: Prentice Hall, 1997. ISBN 0-13-314899-8. Cited 9 times on the pages 3, 6, 7, 16, 18, 19, 20, 21, and 36.
- CALHOUN, P. An entry flight controls analysis for a reusable launch vehicle. In: *Proceedings of the AIAA Guidance, Navigation, and Control Conference*. Denver, CO, USA: American Institute of Aeronautics and Astronautics, 2000. AIAA 2000-1046. Cited on the page 6.
- CRITICISM of the Space Shuttle program. [S.l.]: Wikipedia, 2025. <https://en.wikipedia.org/wiki/Criticism_of_the_Space_Shuttle_program>. Accessed: 2025-06-22. Cited on the page 1.
- DLR-CNES-JAXA. *CALLISTO - Cooperative Action Leading to Launcher Innovation in Stage Toss-back Operations*. 2025. <<https://en.wikipedia.org/wiki/CALLISTO>>. Accessed: 2025-08-13. Cited on the page 2.
- DOYLE, J. Guaranteed margins for lqg regulators. *IEEE Transactions on Automatic Control*, v. 23, n. 4, p. 756–757, 1978. Cited on the page 20.
- DOYLE, J. C.; STEIN, G. Robustness with observers. *IEEE Transactions on Automatic Control*, v. 24, n. 4, p. 607–611, 1979. Cited on the page 20.
- HANSON, J. M. *Advanced Guidance and Control for Reusable Launch Vehicles*. [S.l.], 2000. Disponível em: <<https://ntrs.nasa.gov/api/citations/20000070419/downloads/20000070419.pdf>>. Cited 3 times on the pages 2, 6, and 9.
- KIM, S.; LEE, J.; PARK, H. Linearized state-space model-based attitude control for rocket with four controllable fins – part 1: Basic modeling and identification. *IEEE Access*, IEEE, v. 11, p. 115687–115701, 2023. Disponível em: <<https://doi.org/10.1109/ACCESS.2023.3324114>>. Cited on the page 2.
- MATHWORKS. *MATLAB Control System Toolbox: c2d (Continuous to Discrete) Documentation*. [S.l.], 2024. Compares ZOH, Tustin (bilinear) and other methods; discusses frequency warping and prewarping. Disponível em: <<https://www.mathworks.com/help/control/ref/c2d.html>>. Cited on the page 36.
- MUSGRAVE, J. *Application of LQG/LTR to Reusable Rocket Engine Control*. [S.l.], 1991. Disponível em: <<https://ntrs.nasa.gov/api/citations/19920001910/downloads/19920001910.pdf>>. Cited 3 times on the pages 3, 6, and 17.
- NASA. *Planetary Missions Program Overview*. 2024. <<https://www.nasa.gov/planetarymissions/planetary-mission-program-overview/>>. Accessed: 2025-08-18. Cited on the page 1.

- OLIVEIRA, A. D. *Guidance and Control System Design for Reusable Launch Vehicle Descent and Precise Landing*. Tese (Doutorado) — Politecnico di Milano, Department of Aerospace Science and Technology, Milan, Italy, 2023. Doctoral Dissertation, Cycle XXXVI, Supervisor: Prof. Michèle Lavagna. Cited on the page 1.
- OLIVEIRA, T. D.; LAVAGNA, M. Robust control of reusable launch vehicles in descent phase. *Aerospace*, MDPI, v. 11, n. 11, p. 914, 2024. Disponível em: <<https://www.mdpi.com/2226-4310/11/11/914>>. Cited 10 times on the pages 1, 2, 3, 7, 11, 31, 33, 36, 37, and 38.
- OPPENHEIM, A. V.; SCHAFFER, R. W. *Discrete-Time Signal Processing*. 3rd. ed. Upper Saddle River, NJ: Pearson, 2010. Bilinear transform, frequency warping, and prewarping formulas. Cited on the page 36.
- ORIGIN, B. *NEW GLENN: Our Next (really) Big Step An Orbital Reusable Launch Vehicle that will Build the Road to Space*. 2019. <<https://www.blueorigin.com/new-glenn>>. Accessed: 2025-08-15. Disponível em: <<https://www.blueorigin.com/new-glenn>>. Cited on the page 2.
- POSTLETHWAITE, I. *Multivariable feedback control: analysis and design*. [S.l.]: John Wiley & Sons, Inc., 1996. Cited on the page 23.
- SAGLIANO, M. et al. Unified-loop structured h-infinity control for aerodynamic steering of reusable rockets. *Journal of Guidance, Control, and Dynamics*, v. 46, n. 5, p. 815–837, 2023. Disponível em: <<https://doi.org/10.2514/1.G007077>>. Cited 13 times on the pages 2, 3, 7, 9, 10, 11, 13, 28, 31, 33, 36, 37, and 38.
- SEBORG, D. E.; EDGAR, T. F.; MELLICHAMP, D. A. *Process Dynamics and Control*. 2. ed. Hoboken, NJ, USA: John Wiley & Sons, Inc., 2003. Cited on the page 23.
- SILVEIRA, A. et al. A case study on the attitude control of the vls-1 launcher in the vibratory regime. *IEEE Transactions on Industry Applications*, 2025. Cited on the page 23.
- SpaceX. *Wow! SpaceX Lands Orbital Rocket Successfully in Historic First*. 2015. <<http://www.space.com/31420-spacex-rocket-landing-success.html>>. Accessed: 09/03/2023. Disponível em: <<http://www.space.com/31420-spacex-rocket-landing-success.html>>. Cited on the page 2.
- STEVENS, B.; LEWIS, F.; JOHNSON, E. *Aircraft control and simulation: dynamics, controls design, and autonomous systems*. 3rd. ed. New Jersey: John Wiley and Sons, 2016. Cited 17 times on the pages 3, 4, 6, 13, 14, 16, 17, 18, 19, 20, 21, 22, 23, 33, 36, 37, and 38.
- ZHOU, K.; DOYLE, J. C.; GLOVER, K. *Robust and optimal control*. [S.l.]: Prentice Hall, 1996. Cited on the page 20.