

Madrid, Spain

May 5th-7th

2026

uc3m

Universidad
Carlos III
de Madrid

Robust Attitude Control System Design for the Autonomous Recovery of a Flexible Fairing

Montero Miñán, Alejandro

GNC/AOCS Engineer, Indra-Space , Tres Cantos, Spain. amminan@indra.es

Covasan, Victor

GNC/AOCS Engineer, Indra-Space , Tres Cantos, Spain. vcovasan@indra.es

Ghignoni, Pietro

GNC/AOCS Engineer, Indra-Space , Tres Cantos, Spain. pghignoni@indra.es

Vasconcelos, Jose

Head of GNC/AOCS, Indra-Space , Avionics BU, 1070-061 Lisbon, Portugal jmvasconcelos@indracompany.com

Yábar Valles, Cèlia

GNC Systems Engineer, European Space Agency , Noordwijk, Netherlands celia.yabar.valles@esa.int

ABSTRACT

This paper presents the design of a robust attitude control system for the autonomous recovery of a flexible launch vehicle fairing, a topic currently absent from the literature. The fairing's aeroelastic properties introduce significant control challenges, particularly during atmospheric re-entry. To address these, the study goes beyond controller gain optimization to consider subsystem architecture selection and integration with real actuators, essential for realistic implementation. A linear time-invariant (LTI) model is developed to capture the key dynamics, incorporating parametric uncertainties that reflect aerodynamic variability and structural flexibility. The control architecture is based on \mathcal{H}_∞ synthesis, chosen for its robustness and compatibility with gain-scheduling strategies. A detailed tuning methodology is presented, and the controller is evaluated for both performance and robustness. Validation is performed in a high-fidelity, nonlinear six-degrees-of-freedom simulation environment that includes aeroelastic effects, demonstrating reliable behavior across diverse flight conditions. The resulting controller, combined with simplified navigation and guidance functions, improves recovery success from nearly 0% to 75% under current mission requirements. This approach supports reliable, reusable fairing recovery through advanced guidance, navigation, and control techniques.

Keywords: \mathcal{H}_∞ Control; Attitude Control; Flexible structures; Autonomous recovery; Reusable Fairing

Nomenclature

X_0	=	Any X quantity particularized at the linearization condition.
Δ	=	Variations from the linearization condition.
p, r	=	X and Z comp. of the ang. vel. vector between body and ECI in body coordinates, rad/s.
I_{xx}, I_{zz}	=	Main inertia moment around X and Z body axes, kg·m ² .
I_{xz}	=	Inertia cross product for axes X and Z, kg·m ² .
V	=	Fairing velocity modulus wrt ECI, m/s.



Q	=	Dynamic pressure, Pa.
S	=	Reference surface for the aerodynamic coefficients, m^2
b	=	Reference length for the aerodynamic coefficients, m
α, β, σ	=	Aerodynamic angles (attack, sideslip, bank), rad.
C_{LL}, C_{NL}	=	Aerodynamic roll and yaw moment coefficients reduced at the CoG.
C_S	=	Aerodynamic sideforce coefficient.
m	=	mass, kg.
g	=	Gravity, m/s^2
γ	=	Aerodynamic flight path angle, rad.
L	=	Aerodynamic lift, N.
x_{RCS}	=	RCS cluster location for yaw force, in Body reference frame, X component, m.
x_{cog}	=	Centre of gravity location for yaw force, in Body reference frame, X component, m.
T_y	=	Lateral force provided by the actuators aligned with Body Y axis, N.
T_{max}	=	Maximum thrust provided by the considered RCS, N.
M_x	=	Roll torque provided by the actuators around X body axis, $N \cdot m$.

1 Introduction

Among the various subsystems and components of launch vehicles, payload fairings (PLFs) are unique as they are a multi-functional part with many subsystems, associated costs, and complexity, and yet they are jettisoned within the first few minutes of flight. Despite this high intricacy, value, and their early release along the launch trajectory most launch providers continue to use expendable fairings, with SpaceX being the only exception. Achieving full reusability requires modifications to existing subsystems and the addition of new ones to support extended missions and recovery operations. One such additional subsystem is the guidance, navigation, and control (GNC) system.

Fairing recovery has already been demonstrated successfully by SpaceX for the Falcon 9 launcher [1, 2]. However, no comparable capability currently exists within the European launcher ecosystem. To address this gap, the European Space Agency (ESA) initiated the REFAIR project, awarding Indra-Space the task of developing a GNC subsystem, up to Model-In-the-Loop (MIL) validation level, for the recovery of a 5.4 m class fairing manufactured by Beyond Gravity. A key element in achieving this objective is the design of an attitude control system for the atmospheric re-entry phase, which is the focus of this paper.

Attitude control during re-entry has been a topic of interest for decades, beginning with early crewed capsules such as Apollo, which in turn influenced later designs like the Mars Science Laboratory's attitude control system [3]. This progression continued with the Space Shuttle [4] and extends to more recent ESA re-entry demonstrators, including IXV [5], SpaceRider [6], and EFESTOGNC [7], all of which were led or supported by Indra-Space.

This work presents the controller design for a flexible fairing, covering its architecture selection, analysis, and validation in a high-fidelity environment. It builds on earlier phases of the REFAIR project [8], where two preliminary control architectures, a phase-plane controller and an \mathcal{H}_∞ -based controller, were compared. The phase-plane controller was initially preferred despite certain limitations. Subsequent advances in modelling fidelity and synthesis methods, however, substantially improved the performance of the linear controllers. As will be shown, these refinements motivated the choice of a linear architecture for the final design, which is scheduled for extensive verification and validation in upcoming phases. To provide context, this paper also revisits elements of the earlier work, including the mission overview and the linear time-invariant (LTI) synthesis model, and highlights the improvements introduced since the previous phase.

The remainder of this paper is organized as follows. Section 2 introduces the mission and problem definition, including key structural and aerodynamic characteristics and the concept of operations. Section 3 outlines the development of the LTI model and its associated uncertainties. Section 4 discusses the control system architecture and the rationale behind key design decisions. Section 5 details the controller design process, covering synthesis framework, tuning strategy, and linear analysis results. Section 6 presents the performance of the final design in high-fidelity, nonlinear simulations that incorporate aeroelastic effects. Finally, Section 7 summarizes the main findings and outlines directions for future work.

2 Mission and Problem Description

The GNC subsystem developed during the REFAIR project is designed to enable the recovery of an Ariane 6 payload fairing. To provide context for the controller design presented in this work, the main mission phases are summarized below, highlighting the phase on which this study focuses.

The mission profile, illustrated in Figure 1, is based on preliminary return scenarios analyzed by Indra-Space in earlier studies [9, 10]. It is representative of the expected separation and recovery sequence for next-generation European payload fairings.

The mission starts at fairing separation, at around 70 km altitude and finishes when the fairing lands on top of a recovery barge. The controller design and results presented in this paper focus on the *Re-entry phase*, in which the fairing is in free fall without any braking devices deployed. This phase occurs right after the detumbling and pointing (DET-POINT in Fig. 1), required to reduce the system’s angular velocity and ensure appropriate attitude for the atmospheric part, and the first pilot parachute deployment at approximately Mach 0.7. For more information about the mission, the reader is referred to [9, 10] or the Appendix.

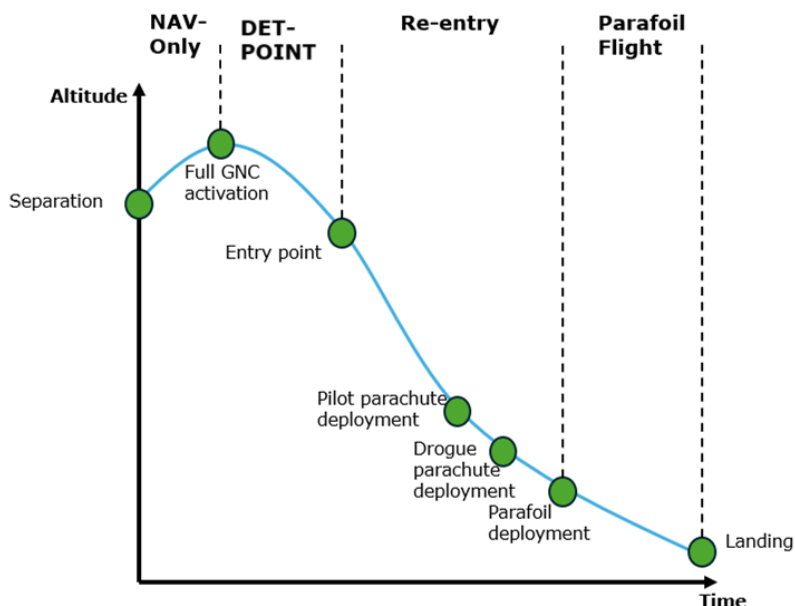


Fig. 1 REFAIR: mission phases

Another crucial aspect for understanding the mission context is a description of the system itself. The GNC subsystem is installed in one half of the fairing, which measures 20 m in length and 5.4 m in diameter. A total of 10 thrusters are mounted on the fairing: six near the nose to provide yaw moments and four closer to the fairing base to provide pitch and roll moments. The number of thrusters was carefully selected so that over 90% of the expected torque commands could be delivered, taking into account variations in trajectory and fairing properties. Their placement and orientation were also optimised to guarantee full control authority over all three body axes while respecting the manufacturer’s prohibited locations and meeting the torque requirements of each channel. Additionally, the thruster positions were chosen to minimise interactions between the actuation and the fairing’s flexible dynamics. A sketch of the fairing, showing its dimensions and eight of the ten thrusters numbered, is presented in Figure 2; the remaining two thrusters are located at the same positions as the central ones but point in the opposite direction.




Fig. 2 Stylised fairing sketch with RCS thruster layout.

In the same fashion as reported in [8], two main characteristics of the system are particularly relevant for the control problem and are described below: aerodynamic static stability and flexibility. The remaining data necessary for the modelling and provided by the fairing manufacturer Beyond Gravity, is highly representative but cannot be disclosed due to confidentiality constraints.

The aerodynamic database for a fairing of this class was built in a previous project [9], providing stability information as well. The longitudinal channel is stable along most of the expected trajectory, up to the transonic regime, where the trim points correspond to statically unstable configurations. The trim angle of attack varies from approximately 50° to 70° as the Mach number decreases. The directional channel, associated with roll, is stable throughout, while the lateral channel exhibits nearly neutral stability in the supersonic regime and becomes unstable between Mach 0.8 and 2.

The current aerodynamic database does not include damping coefficients, which are present in the real system but difficult to characterize at this stage. This omission results in a lack of aerodynamic damping, producing oscillatory behavior around the static stability conditions. For instance, the uncontrolled aerodynamic angle of attack during re-entry of the fairing is illustrated in Figure 3.

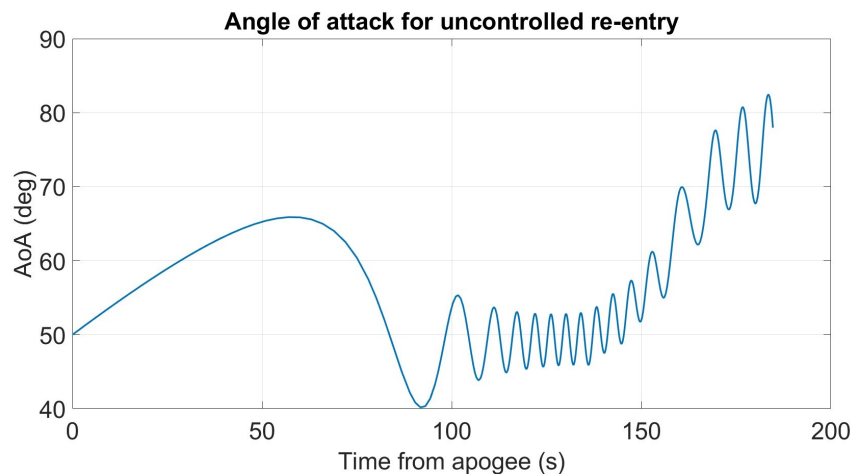


Fig. 3 Uncontrolled angle of attack during re-entry. From [8].

These oscillations are further amplified by the fairing's flexibility. The fairing half is sufficiently compliant to deform during re-entry, altering its shape and, consequently, its aerodynamic characteristics. This generates an oscillatory coupling: aerodynamic loads deform the fairing, which shifts the trim angle of attack to a lower-load configuration, allowing the fairing to return toward its original shape and angle of attack.

After extensive study, it was found that a single mode, termed the "breathing" mode for its characteristic opening and closing motion, accounts for most of the aeroelastic interactions. Its frequency is low enough to fall within the usual control bandwidth. Nevertheless, it was determined that attempting to prevent excitation of this flexible mode during re-entry via control is both ineffective, since aerodynamic forces excite the mode more strongly, and overly restrictive for the control design.

3 Linear Model Description

This section describes the LTI model used to represent the fairing dynamics during the re-entry. Both the nominal model and the uncertainties considered are included in their respective sections.

3.1 Nominal model

As discussed later in Section 4, the control problem focuses primarily on the lateral-directional dynamics. This separation between lateral-directional and longitudinal channels is possible due to the fairing's properties, some of which were simplified to reduce the complexity of the equations, for example, assuming that the center of gravity lies on the fairing centreplane and selecting a reference trajectory with zero sideslip angle ($\beta_0 = 0$). Most of the remaining assumptions required to derive these equations are detailed in [11], which forms the basis for the LTI system in eq. 1. The LTI presented here is identical to the one shown in [8]. Although the underlying dynamics are the same, eq. 1 extends the equations presented in [11], generalizing them for vehicles with significant products of inertia I_{xz} and including the effect of an ideal single RCS cluster on attitude. The output of a controller designed using this LTI model must subsequently be properly allocated to the real thruster cluster.

$$\begin{aligned} \begin{bmatrix} \Delta \dot{p} \\ \Delta \dot{r} \\ \Delta \dot{\beta} \\ \Delta \dot{\sigma} \end{bmatrix} &= \begin{bmatrix} 0 & 0 & \frac{Q_0 S b}{I_{xx} I_{zz} - I_{xz}^2} (I_{zz} \frac{\partial C_{LL}}{\partial \beta} - I_{xz} \frac{\partial C_{NL}}{\partial \beta}) & 0 \\ 0 & 0 & \frac{Q_0 S b}{I_{xx} I_{zz} - I_{xz}^2} (-I_{xz} \frac{\partial C_{LL}}{\partial \beta} + I_{zz} \frac{\partial C_{NL}}{\partial \beta}) & 0 \\ \sin \alpha_0 & -\cos \alpha_0 & \frac{Q_0 S}{m V_0} \frac{\partial C_S}{\partial \beta} & \frac{g \cos \gamma_0}{V_0} \\ \cos \alpha_0 & \sin \alpha_0 & \frac{L_0}{m V_0} - \frac{g}{V_0} \cos \gamma_0 - \frac{\tan \gamma_0 Q_0 S}{m V_0} \frac{\partial C_S}{\partial \beta} & \frac{\tan \gamma_0}{m V_0} L_0 \end{bmatrix} \begin{bmatrix} \Delta p \\ \Delta r \\ \Delta \beta \\ \Delta \sigma \end{bmatrix} \\ &+ \begin{bmatrix} \frac{I_{xz}(x_{RCS} - x_{cog})}{I_{xx} I_{zz} - I_{xz}^2} & \frac{I_{zz}}{I_{xx} I_{zz} - I_{xz}^2} \\ \frac{I_{xz}(x_{RCS} - x_{cog})}{I_{xx} I_{zz} - I_{xz}^2} & -\frac{I_{xz}}{I_{xx} I_{zz} - I_{xz}^2} \\ \frac{1}{m V_0} & 0 \\ \frac{\tan \gamma_0}{m V_0} & 0 \end{bmatrix} \begin{bmatrix} \Delta T_y \\ \Delta M_x \end{bmatrix} \end{aligned} \quad (1)$$

Equation 1 clearly shows the coupling between the β and σ channels, which arises from a nonzero angle of attack and a nonzero inertia product I_{xz} .

Due to the significant variability in the LTI parameters along the reference trajectory, 20 different linearization conditions were selected. This variability motivated the choice of a gain-scheduling control architecture, as discussed in Section 4. The 20 linearization points were carefully chosen to ensure that the system dynamics at each point were sufficiently similar to those of the surrounding points. Figure 4 illustrates this, showing Bode plots of the input-output relationships for the two most relevant channels, β and σ .

An in-depth analysis of the open-loop plants revealed that each linearization point has at least one right-half-plane pole, rendering it unstable without active control. Fortunately, all unstable poles are at low frequency, so the minimum bandwidth required for the controller is not overly restrictive.

These LTI models were validated against numerical linearizations of the high-fidelity 6-DoF simulation model, demonstrating an almost perfect match. The only notable differences occur at very low frequencies, where assumptions such as neglecting Earth's rotation have an impact. An example comparison of analytical and numerical Bode plots for a linearization point at Mach 4.44 is shown in Figure 18 (placed in the Appendix to avoid clutter in this section).

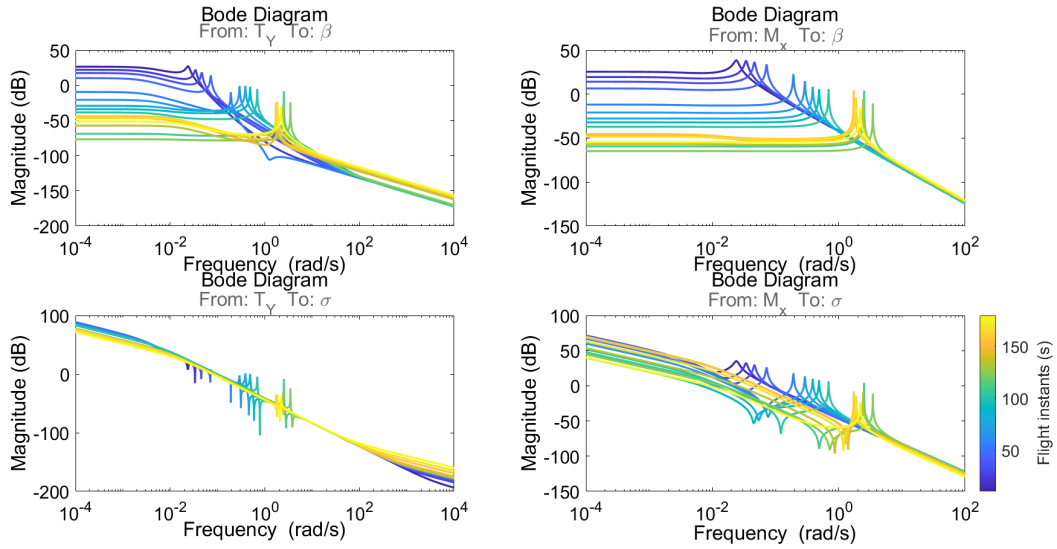


Fig. 4 Input-output bodes for the β and σ channels for the 20 linearization points considered.

The linearization conditions were chosen for symmetric flight ($\beta_0 = \sigma_0 = 0$), zero roll and yaw rates ($p_0 = r_0 = 0$), and an aerodynamic angle of attack following the natural trim dynamics. Given the significant oscillations in this channel, a representative angle of attack was approximated using a third-order polynomial closely matching the static trim, as shown in Figure 5.

This simplification of the trim conditions introduces errors in the linear model. However, as detailed in Section 3.2, the simplified oscillations can be captured by modeling uncertainties in the trim angle of attack α_0 and the aerodynamic coefficient derivatives.

3.2 Uncertainties considered

The linear model presented in [8] did not consider any uncertainties. As a significant improvement for the detailed design phase, parametric uncertainties have been introduced to represent not only intrinsic variations in the system parameters and the associated knowledge errors for the GNC, but also unmodelled dynamics arising from aeroelastic effects in the trim angle of attack.

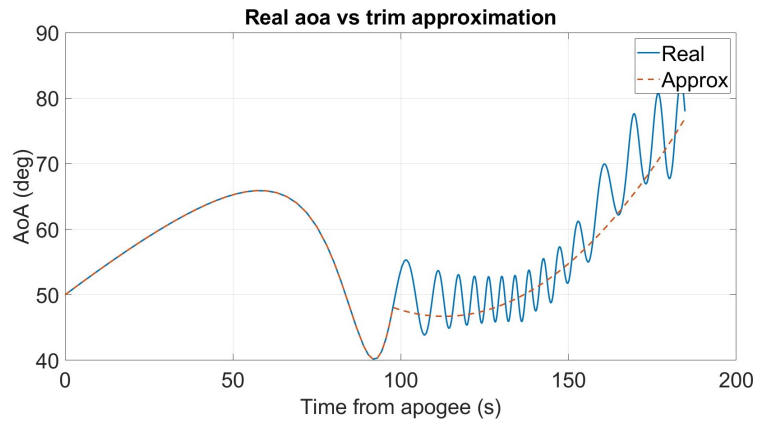


Fig. 5 Trim angle of attack approximation for linearization during re-entry. From [8].

Table 2 lists the uncertainties introduced into the different parameters of the re-entry LTI. Only the cosine of the angle of attack is shown, since the LTI generation script linearly approximates the sine from the cosine through a Taylor expansion to avoid inconsistent sine–cosine combinations.

Table 2 Parametric uncertainties introduced to the LTI

LTI parameter	$\cos \alpha_0$	$\frac{\partial C_S}{\partial \beta}$	$\frac{\partial C_{LL}}{\partial \beta}$	$\frac{\partial C_{NL}}{\partial \beta}$	m	x_{cog}	I_{xx}	I_{zz}	I_{xz}
Uncertainty	$\pm 20\%$	see Fig. 6	see Fig. 6	see Fig. 6	$\pm 5\%$	± 0.1 m	$\pm 5\%$	± 20 kgm ²	$\pm 5\%$

The uncertainties in the Mass, Centre of gravity and Inertia (MCI) properties were derived directly from dispersion data for the fairing. In contrast, the aerodynamic uncertainties, including those associated with the trim angle of attack, combine both the intrinsic aerodynamic variations of the fairing and additional variability induced by oscillations in the trim condition. For each relevant parameter, the mean and standard deviation were computed from the aerodynamic database, evaluated along the reference trajectory with uncertainties enabled.

To capture additional variability from trim changes and the effects of structural flexibility, the total uncertainty used in the LTI generation was conservatively set to a 6σ deviation from the nominal values, rather than the usual 3σ . In this way, flexibility effects are not introduced explicitly into the linear model, but are instead accounted for implicitly through their dominant impact on the aerodynamic coefficients. This avoids unnecessary complication of the model while retaining their influence. The nominal values and 6σ bounds are shown in Figure 6. The uncertainty in $\frac{\partial C_{LL}}{\partial \beta}$, which represents the variation of the roll torque coefficient with the sideslip angle, increases sharply beyond $t \approx 150$ s, corresponding to the transonic phase of the reference trajectory.

This large uncertainty proved problematic for controller synthesis, as no single gain set could stabilize the widely differing LTIs generated under this range. Consequently, the uncertainty in $\frac{\partial C_{LL}}{\partial \beta}$ was capped at 60% of its nominal value, shown as dashed red lines in Figure 6. This value was determined iteratively and should not be interpreted as a strict bound on the actual parametric uncertainty, but rather as a level compatible with the targeted mission success rate. This result highlights the need for significant effort to reduce uncertainty in this parameter during future phases of the REFAIR project, for example, through more detailed CFD studies in the transonic regime or dedicated wind-tunnel campaigns. As a preliminary step, the results from the robust stability and performance analysis could help define the required fidelity for the parameter characterization.

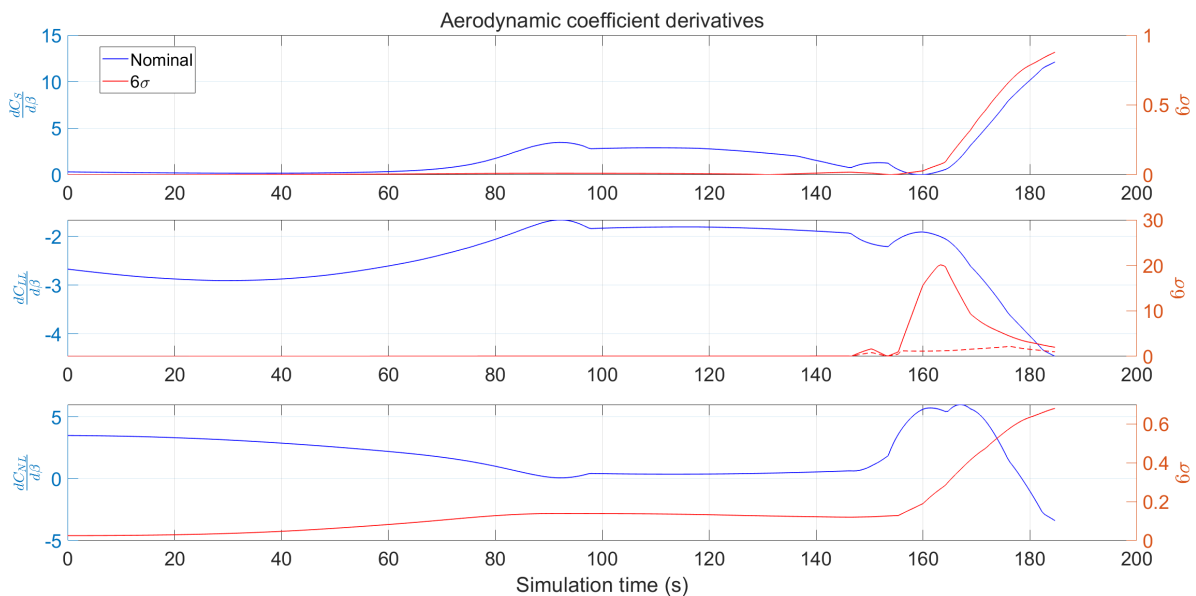


Fig. 6 Uncertainties considered for the aerodynamic coefficients.

4 Architecture selection

Several key aspects must be addressed when defining the controller architecture:

- **Scope of control:** whether to implement full control or only lateral-directional control.
- **Dynamic variability:** how to handle the large variations in system dynamics from a control perspective.

- **Model inputs and allocation function:** which variables should be considered as inputs for the linear model. The allocation function depends on this decision and is described in this section as well.
- **Gain structure:** the type of gain architecture to be used.
- **Reference filtering:** necessity to include or not a reference governor or feedforward term.

4.1 Scope of control

Given the fairing and trajectory characteristics, the longitudinal and lateral–directional channels can be considered decoupled. Moreover, the longitudinal dynamics are aerodynamically stable for most of the trajectory, raising the question of whether controlling the longitudinal channel is necessary. Initially, this channel was left uncontrolled, as noted in [8]. However, this created performance problems: the previous phase’s controller could not always ensure that the angle of attack derivative at the start of the aerodynamic phase matched that of the uncontrolled dynamics (Figure 3), which serves as the reference for linearizing the equations of motion. In some cases, the derivative was positive as required; in others, it was negative, introducing errors into the lateral-directional linear model and, consequently, the control law.

The decision to leave the longitudinal channel uncontrolled was initially motivated by the large torque requirements during the aerodynamic phase of re-entry, which would significantly increase both fuel consumption and the number of thrusters required. However, during the first 30 seconds of this phase, the dynamic pressure is negligible, meaning the thrust and fuel required to control the longitudinal channel are minimal. As a result, a simple longitudinal controller was incorporated into the full control architecture to reduce dispersions in the angle of attack derivative when transitioning from the detumbling and pointing mode. Assuming negligible aerodynamic effects and given the decoupling mentioned above, the angle of attack and pitch angular velocity can be controlled independently using a simple PD controller. The goal is not to track an exact angle of attack profile but to ensure that the uncontrolled evolution during re-entry closely matches the one used for linearisation. The controller gains were tuned iteratively, balancing tracking performance and fuel consumption in simulation.

The final control regulates both longitudinal and lateral-directional channels for the first 30 seconds of the re-entry and only the lateral-directional channels from then.

4.2 Dynamic variability

To address the highly varying dynamics along the trajectory, a gain-scheduling approach [12] is employed, in which the trajectory is linearized at multiple flight conditions, and a dedicated controller is designed for each point. Between these points, the individual controller gains are linearly interpolated. This method provides stability guarantees only for the gains designed at the linearization points, not for the interpolated values. Nevertheless, given the strong track record of success in launcher and re-entry vehicles [5, 13–15], as well as the extensive prior experience within Indra-Space using this architecture, it was chosen for this application.

To improve the gain-scheduling robustness, altitude, instead of time, is selected as scheduling variable. Mach could not be used, as it is not monotonically decreasing during the trajectory.

4.3 Model inputs and allocation function

To simplify the lateral–directional linear model and allow a more flexible allocation scheme, the controller is synthesised using lateral force and roll torque as outputs, as if they were provided by an ideal RCS cluster. Within the controller, the lateral force is then converted into the required yaw torque based on the assumed location of the ideal RCS cluster.

The commanded torques, including pitch during the first 30 seconds, are subsequently allocated to the real RCS cluster through a separate allocation function, briefly described here for completeness. These torques are mapped to individual thruster forces using the nominal RCS locations and the fairing's centre of gravity. As a result, the algorithm is sensitive to errors in RCS position and orientation, as well as to variations in the mass and inertia properties. Any discrepancies between the assumed and actual parameters must therefore be compensated by the controller itself, which should be designed to remain robust against input disturbances.

The individual thruster allocation is divided into yaw and pitch+roll components. The yaw torque can be assigned directly to the corresponding thrusters, independent of the other axes, due to the chosen RCS configuration. The algorithm first determines the sign of the commanded torque. For a positive torque, forces are assigned to THR1 and, if necessary, THR7 and THR9. For a negative torque, the involved thrusters are THR2, THR8, and THR10. The assignment procedure is identical regardless of the torque sign, so it is described here for the positive case only.

Once the sign is determined, the full commanded torque is first assigned to THR1. The required thruster force is computed from the commanded yaw torque $M_{\text{cmd},3}$ as:

$$T_{\text{RCS}} = \min \left(T_{\text{max}}, \frac{M_{\text{cmd},3}}{x_{\text{RCS},1} - x_{\text{COG}}} \right) \quad (2)$$

If this force is below the thruster's maximum, the assignment is complete. Otherwise, additional thrusters are activated to provide the remaining torque. The force for the extra thruster (THR7) is calculated as:

$$T_{\text{RCS},7} = \min \left(T_{\text{max}}, \frac{M_{\text{cmd},3} - T_{\text{max}}(x_{\text{RCS},1} - x_{\text{COG}})}{x_{\text{RCS},7} - x_{\text{COG}}} \right) \quad (3)$$

Similarly, the next thruster (THR9) is used if more force is needed:

$$T_{\text{RCS},9} = \min \left(T_{\text{max}}, \frac{M_{\text{cmd},3} - T_{\text{max}}(x_{\text{RCS},1} - x_{\text{COG}}) - T_{\text{max}}(x_{\text{RCS},7} - x_{\text{COG}})}{x_{\text{RCS},9} - x_{\text{COG}}} \right) \quad (4)$$

If the requested torque still exceeds the maximum combined force, the actuator suite saturates, providing only the resultant torque from full-force commands to THR1, THR7, and THR9 (or THR2, THR8, and THR10 for negative torque). The actuators were sized to cover over 90% of the expected commanded torques; designing for full 100% coverage would have required a significantly larger number of thrusters and fuel. This compromise aligns with the required reliability and performance of the GNC system.

Roll and pitch allocations must account for two critical factors. First, yaw thrusters generate an undesired roll torque due to the difference in Z coordinate between their position and the centre of gravity, which must be compensated. Second, roll and pitch thrusters are shared between both axes, so their allocation should be performed jointly. To address the first point, an estimate of the undesired roll torque produced by the yaw thrusters is computed and subtracted from the commanded roll torque. For a positive yaw torque command, this is expressed as:

$$M_{\text{roll,extra}} = -T_{\text{RCS},1}(z_{\text{RCS},1} - z_{\text{COG}}) - T_{\text{RCS},7}(z_{\text{RCS},7} - z_{\text{COG}}) - T_{\text{RCS},9}(z_{\text{RCS},9} - z_{\text{COG}}) \quad (5)$$

The equivalent expression for a negative yaw torque command requires only a sign change in the terms and an update of the relevant RCS indices.

Once this correction is applied, the individual thruster forces can be computed. In the current RCS configuration, thrusters 3 and 5 share the same position, as do thrusters 4 and 6. This is exploited in the actuator management logic by computing the required forces for thrusters 3 and 6 and adjusting the allocation if any resulting force is negative. While a similar approach could be applied when thrusters do not share positions, it would require an additional step, which can be omitted in the present algorithm.

The algorithm initially assumes that the commanded torques can be achieved using thrusters 3 and 6. It then solves the following linear system to determine the individual forces:

$$\begin{bmatrix} M_{\text{cmd},1} - M_{\text{roll,extra}} \\ M_{\text{cmd},2} \end{bmatrix} = \begin{bmatrix} -(y_{\text{RCS},3} - y_{\text{COG}}) & y_{\text{RCS},6} - y_{\text{COG}} \\ -(x_{\text{COG}} - x_{\text{RCS},3}) & -(x_{\text{COG}} - x_{\text{RCS},6}) \end{bmatrix} \begin{bmatrix} T_{\text{RCS},3} \\ T_{\text{RCS},6} \end{bmatrix} \quad (6)$$

If the computed force for THR3 is negative, the same magnitude is assigned positively to THR4, and THR3 is not commanded. The same logic is applied to the THR5 and THR6 pair. This direct reassignment is possible due to the aforementioned shared positioning, but the algorithm can be generalised by resolving the linear allocation problem using the relevant RCS geometry once the active thrusters have been identified.

Given the current stage of the project, no reconfiguration capability has been implemented yet. Nonetheless, the algorithm could be readily extended to incorporate thruster health information and selectively disable or prioritize actuators based on their condition.

4.4 Gain structure

The gain structure adopted to control the lateral-directional system has been derived from previous Indra-Space projects, such as IXV [5] and EFESTOGNC [7].

For each linearization point, the controller architecture is relatively simple: the roll rate p , yaw rate r , and sideslip β channels are regulated through single gains, while the bank angle σ channel is controlled using a PI regulator. The inclusion of the integral term improves the tracking performance of the bank angle, which is the most relevant state for accurate trajectory tracking.

In addition, a high-pass filter is applied in the sideslip channel to prevent continuous actuator activity and to allow the system to stabilise at the trim sideslip angle, even when it differs from zero. Ideally, this term would have been implemented as a simple derivative, but an additional low-frequency pole and a limitation to high-frequency amplification were introduced to avoid implementation issues. The HPF is composed by a low frequency zero, at $\omega = 10^{-4}$ rad and a pole at $\omega = 10$ rad.

A controller sketch for a single output is shown in Figure 7. The same structure is employed for both output channels, yaw lateral force and roll torque, with each channel using different gains.

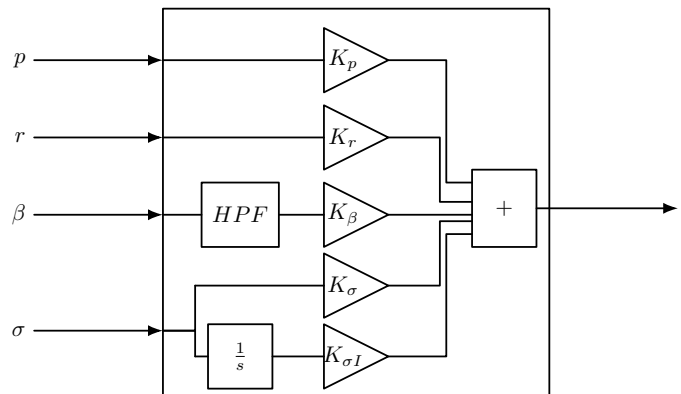


Fig. 7 Controller sketch.

4.5 Reference filtering

It is important to note that, given the assumptions and linearisation conditions applied, these models are in principle valid only around $\beta = \sigma = p = r = 0$. Consequently, controllers derived from them cannot directly handle large or rapid bank commands. Nevertheless, extensive testing, both within and

beyond this project, has shown that the closed-loop system, combining the real (nonlinear) dynamics with controllers based on the linear approximation, can follow large bank commands provided they evolve sufficiently slowly. To enforce this limitation, a reference governor is introduced between the guidance module of the GNC subsystem and the controller.

The reference governor receives the desired bank command from the guidance function and filters it by limiting its rate of change. In addition to ensuring proper tracking performance from the controller, this filtering process also mitigates actuator saturation.

The reference governor is composed by a second order filter with magnitude and rate saturations, as shown in eq. 7, where ω_n and ξ are the natural frequency and damping ratio describing the filter behaviour and $\text{sat}_R(*)$ and $\text{sat}_M(*)$ represent rate and magnitude saturation functions.

$$\ddot{\sigma} = \text{sat}_R \left[\text{sat}_M \left(-2\xi\omega_n\dot{\sigma} + \omega_n^2 \text{sat}_M(\text{sat}_M(\sigma_{\text{ref}}) - \sigma) \right) \right] \quad (7)$$

The final control command is obtained by integrating twice the reference governor output, $\ddot{\sigma}$. Its parameters were designed alongside the thruster selection process to ensure minimal saturation for the expected commands. In particular, the parameter values are shown in table 3.

Table 3 Reference governor parameters

Parameter	ω_n	ξ	Bank saturation	Bank rate saturation	Bank acceleration saturation	Bank acceleration rate sat.
Value	0.19 Hz	0.8	± 90 deg	± 6 deg/s	± 9 deg/s ²	± 7.5 deg/s ³

5 Controller design and analysis

5.1 Synthesis framework and tuning

The gains at each linearisation point are computed using structured H_∞ methods. In particular, the synthesis framework illustrated in Figure 8 is combined with the predefined architecture described in Section 4 and the MATLAB function `systeme`. This function allows the user to explicitly define performance and robustness constraints, such as reference tracking, disturbance rejection, or control effort, which are internally converted into input/output weighting functions for the structured H_∞ optimisation. Additionally, `systeme` enables channel-specific constraints, providing more flexibility than its predecessor `hinfstruct`, where such individual channel specification is less straightforward.

The considered constraints can be divided into hard and soft, depending on whether the algorithm includes them directly into the optimisation problem constraints or in the objective function, respectively. The particular constraint used to obtain the gains for the REFAIR project are:

- Hard constraints:
 - Tracking constraint between the bank reference and the bank signal. A first order transfer function with a time constant of 0.25 s.
 - Rejection constraint between the bank reference and the sideslip channels. This avoids exciting the sideslip channel when a new bank command is sent. A threshold of 0.4 is imposed on the closed-loop gain between the bank reference and the sideslip signal.
 - Input margin constraint. A minimum of 4.5 dB and 20 degs. is imposed at the plant input. Since these correspond to the multiple-loop input-output disk margins from MATLAB, the associated classical margins for each channel will be higher and closer to the industry standard of 6 dB and 30 deg, as shown later in section 5.2.

- Output margin constraint. A minimum of 1.5 dB and 15 degs is imposed at the plant output. Lower values are needed since the noise associated with the navigation function is considered to be lower than the expected input perturbations.
- Minimum pole decay of 10^{-4} . Given that some channels remain unconstrained with the previously defined constraints (apart from stability from the input and output margins), a minimum decay for all the closed loop poles is imposed to avoid significant oscillations.
- Soft constraints: The only soft constraint considered is on the gains from the bank command to the control efforts, both T_y and M_x . By including this constraint in the soft set, the algorithm minimises these gains as much as possible while still satisfying the other constraints, thereby reducing the control effort required to meet all remaining performance requirements.

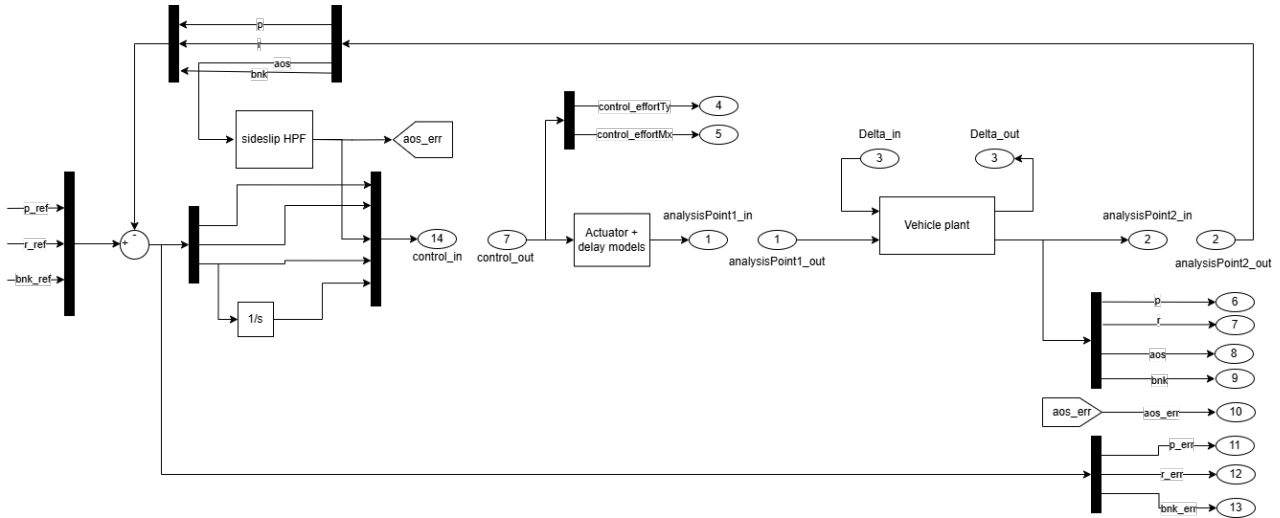


Fig. 8 Synthesis framework sketch from Simulink implementation.

5.2 Synthesis results and controller analysis

The controller gains obtained from the synthesis with the constraints described in Section 5.1 are shown in Figure 9. These gains are then used to assess the closed-loop response of the controlled linear system. The analysis includes verification of the most relevant closed-loop transfer functions, step response evaluations, and stability margin computations.

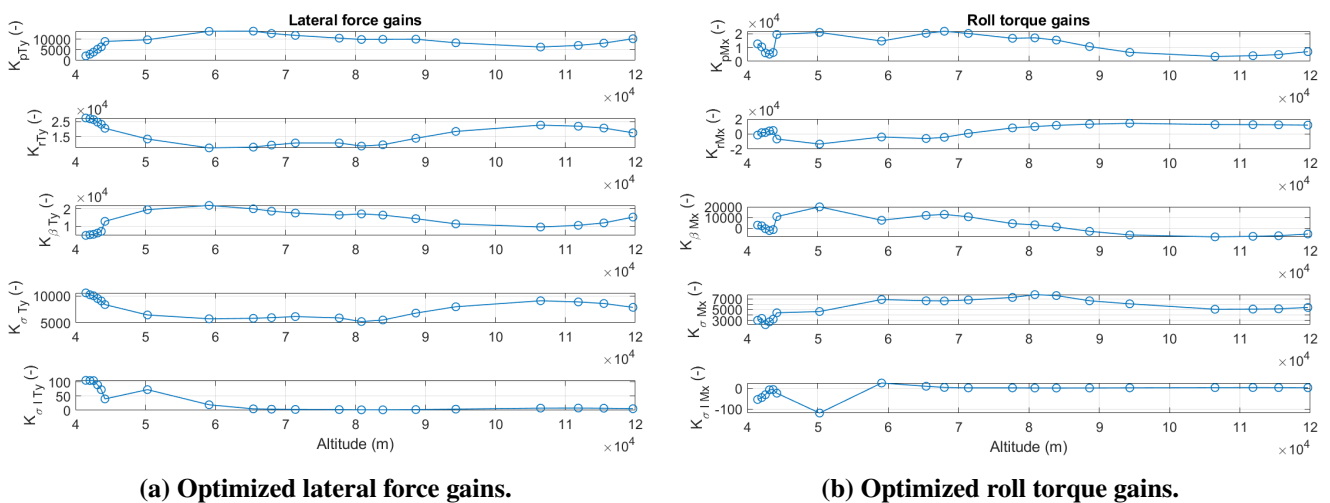


Fig. 9 Resulting gains for the REFAIR project after H_∞ synthesis.

A particularly important case is the step response to a unit bank reference input in the bank and sideslip channels. An example for the 11th linearisation point is shown in Figure 10, chosen as a representative mid-point, given the similarities of the input-output relationship between all the linearization points, as seen in Figure 4. The results indicate that the bank channel reaches the unit reference in about 10 seconds, which coincides with the time required for the sideslip to return to its steady-state value, without exceeding the imposed 0.4 threshold. This behaviour remains consistent even under uncertainty, as demonstrated by the close overlap of the different response curves in the figure.




Fig. 10 Sideslip and bank response to unitary bank reference command.

After verifying the frequency- and time-domain responses, gain and phase margins were evaluated at all linearisation points. For the MIMO system, input margins were computed for each disturbance channel, output margins for each noise channel, and an overall MIMO margin quantified robustness to simultaneous input–output disturbances. Since only the input margins are subject to explicit requirements, they are the ones reported in Figure 11, where channels 1 and 2 correspond to lateral force and roll torque disturbances, respectively. As shown in the figure, all reported margins exceed the specified requirements, confirming the robustness of the closed-loop system. All margins are included in the appendix, in Figure 19.

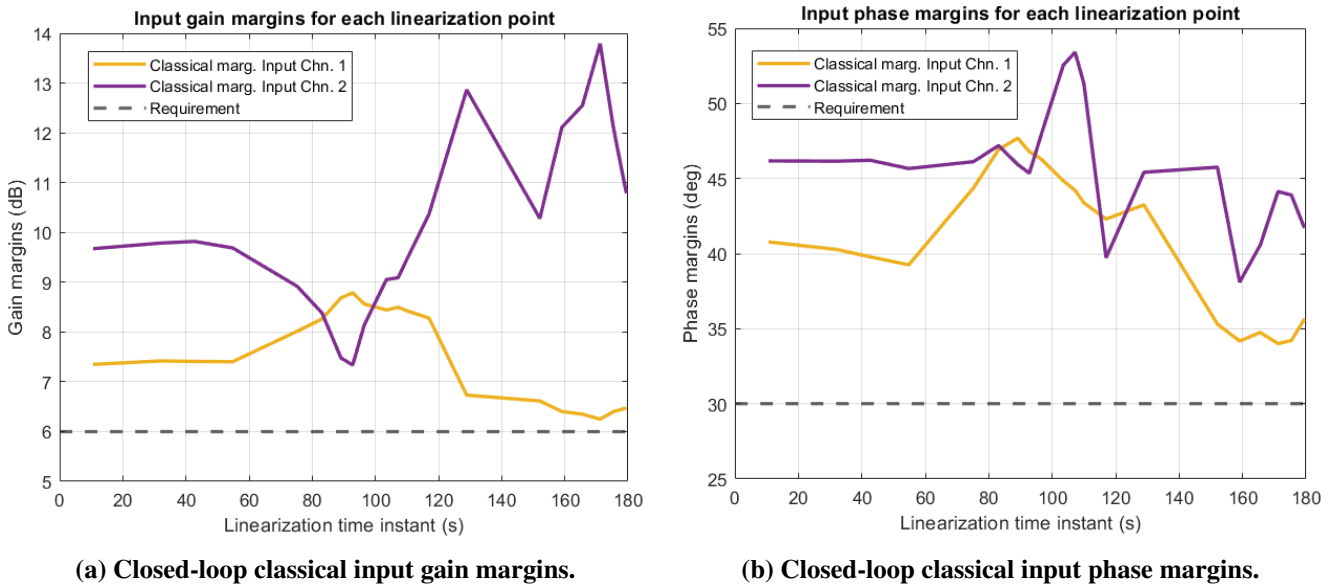
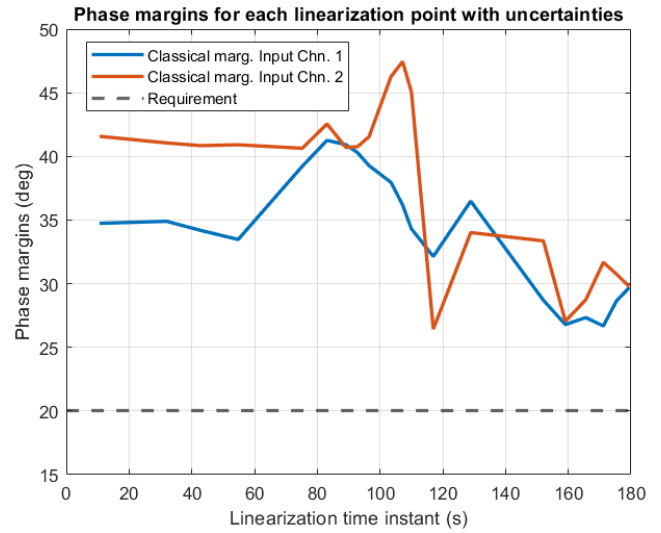
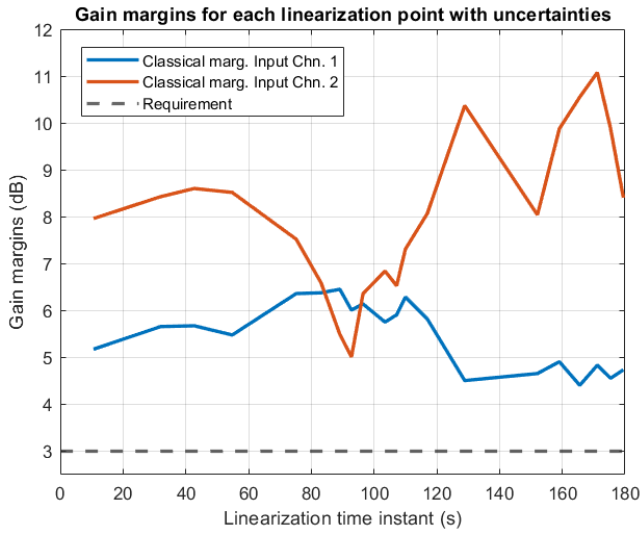


Fig. 11 Closed-loop classical input margins for all linearization points.


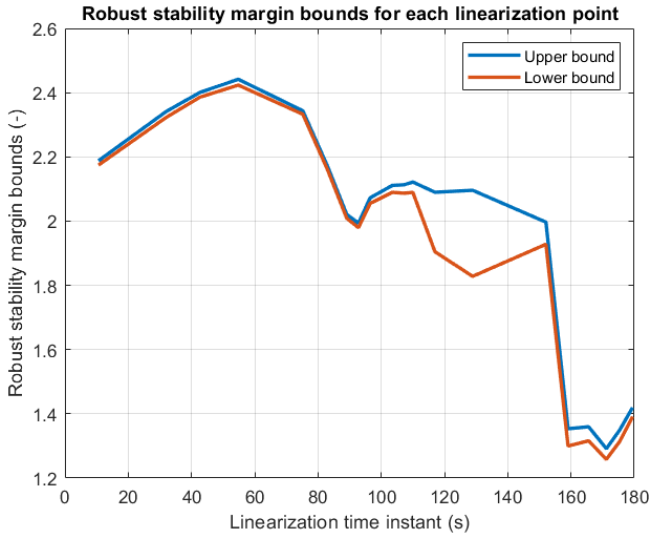
To further assess robustness beyond the nominal plant, gain and phase margins were also evaluated for 200 uncertain plant realisations. For each case, the closed-loop margins were computed, and the minimum values across all realisations were compared against the corresponding requirements, as shown in Figure 12.



(a) Worst-case closed-loop classical input gain margins. (b) Worst-case closed-loop classical input phase margins.

Fig. 12 Worst-case closed-loop classical input margins for all linearization points with uncertainties.

To conclude the closed loop analysis in linear domain, robust stability margins were assessed over the uncertainty range considered, demonstrating that the system remains stable for all cases. The uncertainty in $\frac{\partial C_{LL}}{\partial \beta}$ was limited during synthesis, as the full range proved too large, yet it remains the dominant constraint during transonic flight. This is evident from Figure 13a, compared with Figure 13b where $\frac{\partial C_{LL}}{\partial \beta}$ was fixed at its nominal value: the robust stability margin nearly doubles at the final operating points, where transonic effects amplify this uncertainty. These results confirm the critical impact of $\frac{\partial C_{LL}}{\partial \beta}$ and emphasise the need for improved characterisation in future work.



(a) Robust stability margins with all uncertainties. (b) Robust stability margins without $\frac{\partial C_{LL}}{\partial \beta}$ uncertainties.

Fig. 13 Robust stability margins.

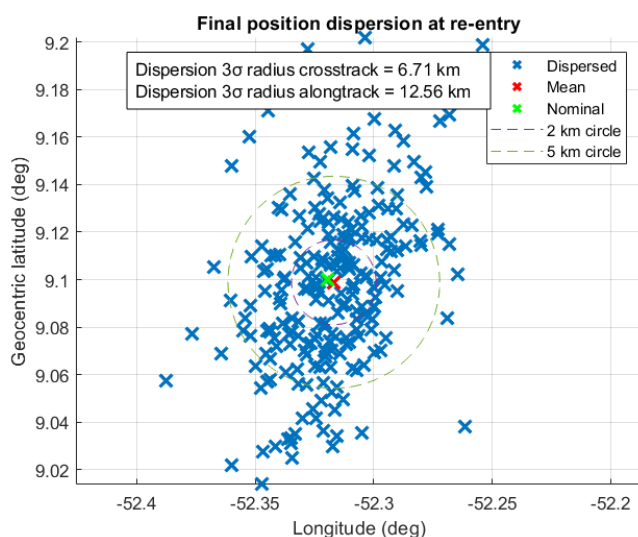
6 Simulation Results

The proposed controller was evaluated in a high-fidelity re-entry simulator, which includes rigid-body dynamics, structural flexibility and its aerodynamic coupling, as well as detailed sensor and actuator models. To isolate control performance, the navigation function was replaced with a performance model


representative of expected estimation errors, and the guidance module with a simple open-loop look-up-table commanding the same bank profile across all runs.

A Monte Carlo campaign of 251 simulations was carried out from separation to pilot-parachute deployment. Variations included MCI properties, flexible modes, sensor and actuator characteristics, atmospheric conditions, and fairing release parameters. To simplify the analysis and isolate the re-entry controller as much as possible, the first two phases from Figure 1 and the re-entry phase were simulated separately, with dispersion results from one phase used to initialise the next. The results presented here cover only the re-entry phase.

Of these runs, 240 remained stable, where stability was defined as the ability to track bank commands without uncontrolled tumbling. Thus, the simplified GNC subsystem achieved stability in over 95% of the cases within the considered uncertainty set. Most instabilities were linked to actuator saturation, consistent with expectations from the design phase. The increase in weight to cover all the torque command expected range was considered prohibitive, even with the controller optimisation process described in section 5.1. The following results focus on the 240 stable cases.



(a) Position dispersion at pilot parachute deployment.



(b) Latitude-longitude maps for the simulated trajectories.

Fig. 14 Trajectory results from the Montecarlo campaign.

A key performance metric is the position dispersion at pilot-parachute deployment (Figure 14a). The simplified subsystem does not meet the 5 km requirement, with along-track dispersion being the dominant contributor. This outcome reflects the limitations of the open-loop guidance algorithm, which cannot compensate for MCI or aerodynamic dispersions. Performance is expected to improve once the presented controller is coupled with a closed-loop guidance law. Nevertheless, the recovery success rate increases significantly: from nearly 0% without active control to 75% when considering landings within the 5km radius, and over 95% when considering all stable cases.

Figure 14b shows representative latitude–longitude trajectories, assuming launch from Kourou with fairing separation at 70 km altitude near 7°N, 52°W. The commanded bank initially drives the fairing northwest, followed by a sign reversal that redirects the trajectory northeastward.

Additionally, fuel consumption and RCS firings were assessed as cost-related figures of merit (Figure 15 and Figure 16). Across the campaign, total fuel consumption remained below the 50 kg limit, with a maximum of 34.42 kg, an average of 12.54 kg, and a 3σ value of 19.65 kg for the re-entry phase. These values exclude the detumbling and pointing phases, which themselves remained below 10 kg, with the number of thruster firings during detumbling always under 30. This confirms that the subsystem operates well within fuel constraints.

Similarly, Figure 16 shows that the total number of firings remains below the maximum of 22,500 per thruster per flight, even when including those from the earlier phases. Another relevant feature in Figure 16 is the distribution of firings among the yaw thrusters (THR1–2 and THR7–10). The higher count for THR1 and THR2, followed by THR7 and THR8, indicates that in most cases two thruster pairs were sufficient to generate the required yaw torques, with the second pair engaged only occasionally. However, there are cases where the torque demand exceeds the capability of two pairs, requiring the third to avoid saturation.

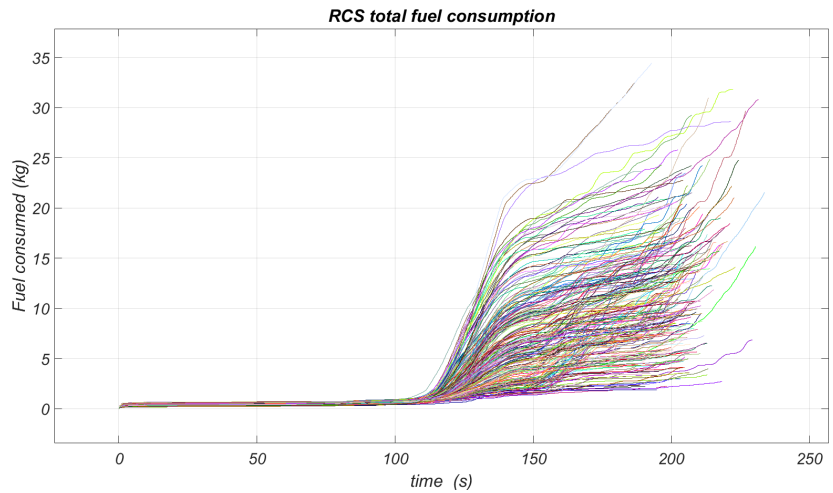


Fig. 15 Fuel consumed before pilot parachute deployment.

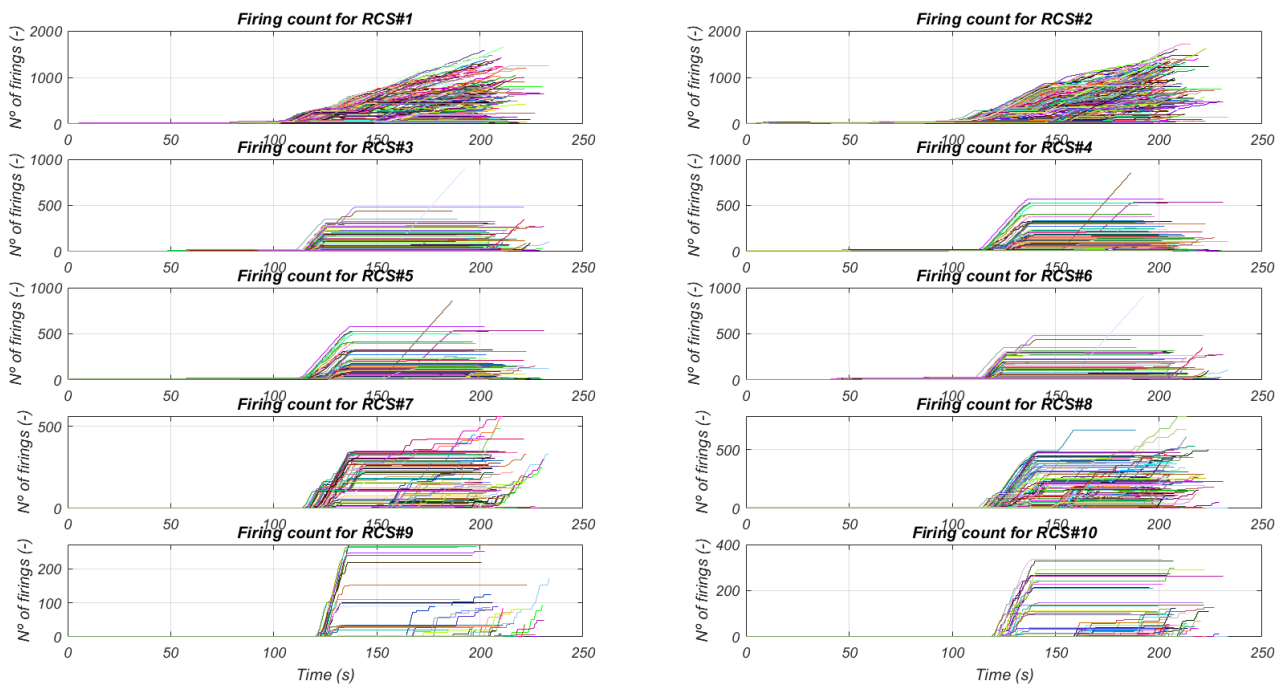


Fig. 16 Synthesis framework sketch from Simulink implementation.

Finally, Figure 17 illustrates the aerodynamic angles during the re-entry phase. The angle of attack closely follows the natural trim dynamics shown in Figure 3, thanks to the longitudinal control applied during the first 30 seconds. The sideslip angle remains bounded, though it does not maintain a fixed reference at zero, to accommodate trim variations caused by dispersions. The bank angle follows a profile starting from -32° , with a bank reversal around 125 seconds and a linear reduction from 200 seconds onwards, consistent with the trajectories shown in Figure 14b. Overall, the controller successfully tracks the bank commands, with only occasional oscillations following the bank reversal. These oscillations are currently hypothesized to result from actuator saturation in specific simulations, without leading to a fully uncontrollable state.

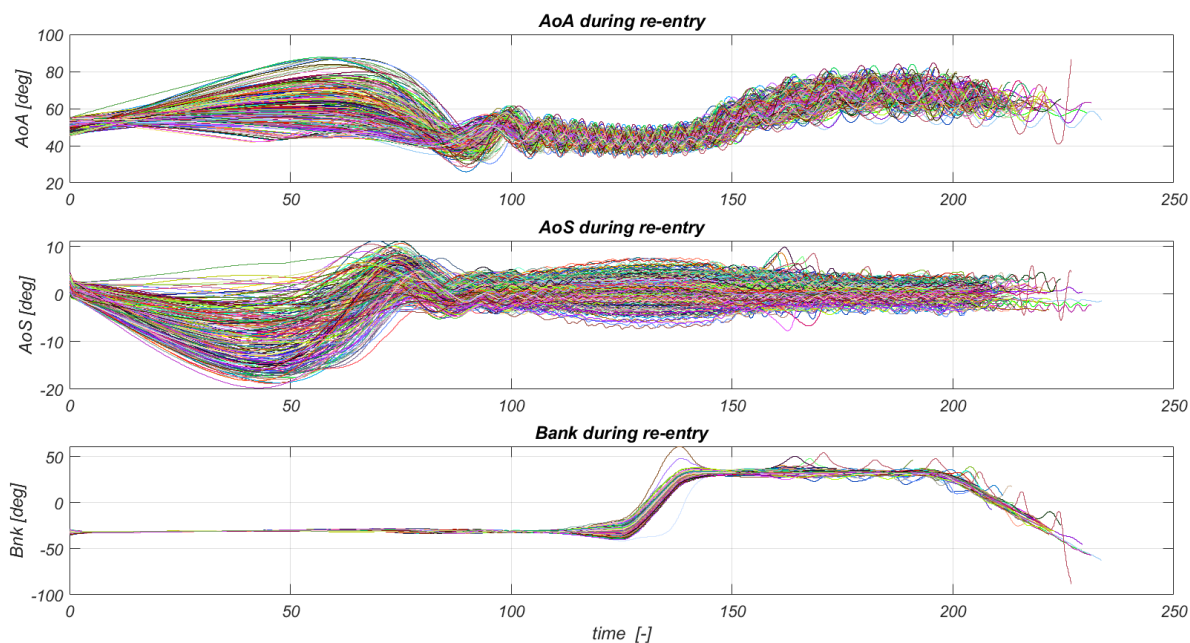


Fig. 17 Aerodynamic angles during re-entry phase for stable simulations.

7 Conclusion

This paper has presented the design process and results of an H_∞ -based robust controller for the recovery of a flexible fairing. When combined with basic navigation and guidance functions, the controller increases the mission success rate from nearly 0% to 95% in terms of stability, and to 75% when performance metrics are also considered, despite fairing parametric uncertainties and variations in initial and atmospheric conditions.

Further refinements of the actuator and controller design are possible. In particular, modifying the actuator suite and/or the reference governor could eliminate saturations, thereby increasing the stability rate to nearly 100% and improving the performance of cases where saturations occurred without causing instability. However, such modifications would impact the mass budget, making it necessary to assess the trade-off between GNC subsystem requirements and overall system constraints.

Finally, landing precision is expected to improve substantially once the controller is paired with a closed-loop guidance algorithm, which should significantly reduce along-track errors compared with the open-loop bank profile results presented here.

Appendix

This appendix provides additional mission details and supplementary images that enhance the reader's understanding, though they are not essential to the core content of the paper.

As mentioned in section 2, the mission begins at fairing separation, typically occurring at altitudes above the Kármán line (100 km). Immediately after separation and jettisoning, all recovery subsystems remain inactive except for the navigation function, which begins computing its state estimates. Initialization is based either on estimates carried over from the launch phase or on data received from the launcher's navigation system. This period, referred to as the *NAV-Only phase*, may include passage through the rocket plume, during which the fairing can experience significant angular rates.

Once the inhibition time elapses, the *Detumbling and pointing phase* begins. The remainder of the recovery system is activated, and the GNC commands onboard RCS thrusters to reduce angular rates.

Following detumbling, the fairing is reoriented to the desired re-entry attitude, so that the atmospheric phase can start with the appropriate initial conditions.

Following this, the mission enters the *Re-entry phase*. As the fairing descends into denser atmosphere, aerodynamic drag progressively slows it. Throughout this phase, the GNC maintains the commanded attitude, ensuring controlled deceleration while avoiding excessive aerodynamic loading or heating. At approximately Mach 0.7, a pilot parachute is deployed, followed four seconds later by a drogue parachute. Although these events occur before the end of the re-entry phase, active guidance and control for re-entry are no longer applied after parachute deployment.

Finally, when the fairing reaches the required Mach and altitude conditions (around Mach 0.2 and 6 km altitude), the parafoil is deployed, initiating the *Parafoil Flight phase*.




Fig. 18 Input-output bode comparison between analytical and numerical models.

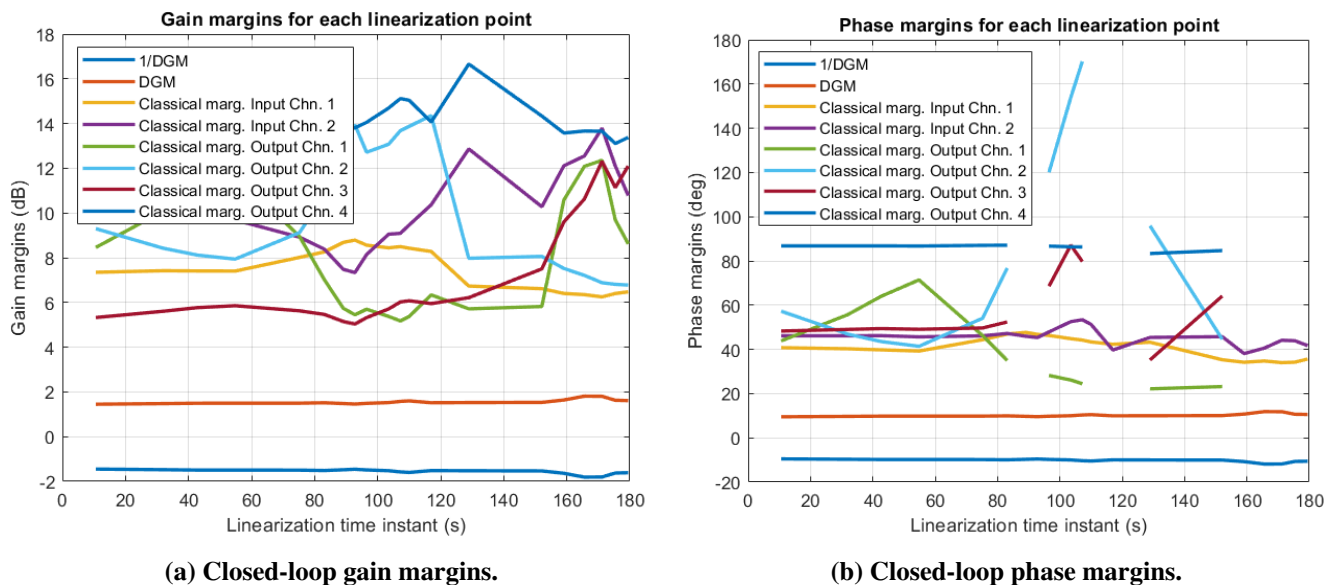


Fig. 19 Closed-loop margins for all linearization points.

In Figure 19, output channels 1–4 correspond to roll rate, yaw rate, sideslip angle, and bank angle noises, while the MIMO margins are labelled as DGM. Note that phase margins are not always finite

in closed-loop MIMO systems; this explains the discontinuous appearance of the output phase margin curves in Figure 19b.

Acknowledgments

The results presented in this paper have been achieved under funding by the Technology Development programme with contract 4000143244/23/NL/CRS. The view expressed in this paper can in no way be taken to reflect the official opinion of the European Space Agency. The authors wish to thank Beyond Gravity and ESA teams for their support to the REFAIR activities.

Declaration of Use of Artificial Intelligence

Artificial Intelligence was not used within the scope of the REFAIR project. However, it was employed for proofreading and for gathering some references during the preparation of this paper. All outputs produced by the AI were carefully reviewed by the authors to ensure the text preserved its original meaning and that the references were relevant.

References

- [1] The Verge. SpaceX attempts to catch falcon 9 fairings with thrusters and parachutes, 2017. Accessed: 2025-08-12. <https://www.theverge.com/2017/3/30/15132314/spacex-launch-fairing-landing-falcon-9-thruster-parachutes>.
- [2] Darrell Etherington. SpaceX launches re-flown fairing for the first time and breaks a falcon 9 booster re-use record, 2019. Accessed: 2025-08-12. <https://techcrunch.com/2019/11/11/spacex-launches-re-flown-fairing-for-the-first-time-and-breaks-a-falcon-9-booster-re-use-record/>.
- [3] P. B. Brugarolas, A. M. San Martin, and E. C. Wong. Attitude controller for the atmospheric entry of the mars science laboratory. In *AIAA Guidance, Navigation and Control Conference and Exhibit*, USA, Aug. 2008.
- [4] Charles Stark Draper Laboratory, Inc. Integrated digital flight control system for the space shuttle orbiter. Technical Report E-2736, Revision 1, Charles Stark Draper Laboratory, Inc., July 1973. Accessed: 2025-08-12. <https://ntrs.nasa.gov/api/citations/19730022103/downloads/19730022103.pdf>.
- [5] D. Gherardi et al. Ixv gnc subsystem design and performances. In *8th International ESA GNC Conference*, Karlovy Vary, Czech Republic, June 2011.
- [6] C. Recupero, G. De Zaiacomo, and A. Tarabic. Space rider: Entry and taem g&c of the future european reusable space transportation system. In *2nd International Conference on High-Speed Vehicle Science Technology (HiSST)*, Bruges, Belgium, Sept. 2022.
- [7] N. Somma, P. Guerreiro, B. Midões, J.P. Belfo, L. García Basabe, A. Princi, G. Guidotti, and J. Vasconcelos. Advanced gnc technology for inflatable heat shields: Preliminary design. In *22nd European Control Conferences (ECC)*, Stockholm, Sweden, June 2024.
- [8] A. Montero Miñán, V. Covasan, A. Tarabic, A. Princi, J. Vasconcelos, P. Rosa, G. De Zaiacomo, and C. Yábar Vålles. Preliminary control architecture selection for fairing atmospheric re-entry problem. In *ROCOND 2025*, Porto, Portugal, July 2025. Presented at the conference, manuscript not yet published online.
- [9] D. Bonetti et al. Reusable payload fairings: Mission engineering and gnc challenges. In *Proceedings of the 7th European Conference for Aeronautics and Space Sciences (EUCASS)*, 2019. EUCASS 2019. doi: [10.13009/EUCASS2019-638](https://doi.org/10.13009/EUCASS2019-638).



- [10] Andrei Filip Cojocaru, Alina Ionita, Andrei Eugen Tarabic, Andrei Lucian Alexe, Ionut Grozea, Paulo Rosa, and Cèlia Yábar Valles. Parafoil mpc guidance and path following control for d&l phase of reusable space vehicles. In *Proceedings of the 73rd International Astronautical Congress (IAC)*, 2022. Paper D1/3. <https://dl.iafastro.directory/event/IAC-2022/paper/68368/>.
- [11] E. Mooij. *Linear Quadratic Regulator Design for an Unpowered, Winged Re-entry Vehicle*. Delft University Press, Netherlands, 1998.
- [12] Wilson J. Rugh and Jeff S. Shamma. Research on gain scheduling. *Automatica*, 36(10):1401–1425, 2000. doi: [10.1016/S0005-1098\(00\)00058-3](https://doi.org/10.1016/S0005-1098(00)00058-3).
- [13] Jay Brandon, Stephen Derry, Eugene Heim, Richard Hueschen, and Barton Bacon. Ares i-x stability and control flight test: Analysis and plans. In *AIAA Space 2008 Conference and Exposition*, San Diego, CA, 2008. doi: [10.2514/6.2008-7807](https://doi.org/10.2514/6.2008-7807).
- [14] John A. Edinger, Thomas W. Hilgeman, and Donald F. Synder. The space shuttle ascent flight control system. In *AIAA Guidance and Control Conference*, San Diego, CA, 1976.
- [15] Diego Navarro-Tapia, Andres Marcos, Samir Bennani, and Christophe Roux. Structured h-infinity control design for the vega vehicle: Recovery of the legacy control behaviour. In *International ESA conference on guidance, navigation and control systems*, Salzburg, Austria, 2017. Discusses gain scheduling in operational Vega launcher TVC. doi: [10.13009/EUCASS2017-257](https://doi.org/10.13009/EUCASS2017-257).