Peer reviewed version

Link to published version (if available):
10.1002/rnc.4109
10.1002/rnc.4109

Link to publication record in Explore Bristol Research
PDF-document

University of Bristol - Explore Bristol Research
General rights

Synthesis and Analysis of Robust Control Compensators for Space Descent & Landing

P. Simplício[1*], A. Marcos[1], E. Joffre[2], M. Zamaro[2] and N. Silva[2]

[1] *University of Bristol, Technology for AeroSpace Control (TASC) Group, Department of Aerospace Engineering, University Walk, Bristol BS8 1TR, UK*
[2] *Airbus Defence and Space Ltd, Gunnels Wood Road, Stevenage SG1 2AS, UK*

SUMMARY

In this article, a complete modelling, synthesis and analysis methodology of control compensators for descent and landing (D&L) on small planetary bodies is presented. These missions are scientifically very rewarding, but technically extremely challenging due to the complex and poorly-known environment around those bodies, which calls for the ability to manage competing robustness and performance requirements. While this issue is typically addressed via the redefinition of D&L guidance strategies, here it is tackled through the augmentation with a simple yet robust control compensator. This compensator is designed using linear fractional transformation (LFT) modelling to capture the interplay with uncertain gravity fields and the recently-developed structured $\mathcal{H}_\infty$ optimisation framework, which has been proved particularly suitable for industry-oriented applications. The proposed approach is completely generic, but employs the scenario of a landing on the Martian moon Phobos as an illustrative example. Different compensators are then verified and compared analytically via the structured singular value μ and through high-fidelity Monte-Carlo simulation. Copyright © 2017 John Wiley & Sons, Ltd.

1. INTRODUCTION

Renewed scientific interest has been growing in the exploration of small asteroids, in addition to larger planetary bodies such as Mars, since their weaker gravity field makes them more easily accessible. However, these missions are very challenging from an engineering point of view, particularly if the natural dynamics in the vicinity of the target asteroid is exploited (to alleviate tight descent & landing (D&L) propellant consumption requirements). This is because small planetary bodies are typically characterised by highly irregular and poorly-known shapes, which render their physical

---

*Prepared using **rncauth.cls** [Version: 2010/03/27 v2.00]*

environment extremely uncertain and variable. Moreover, due to the interplanetary distances involved, fully autonomous guidance and control algorithms are required to cope with communication delays and spacecraft subsystem degradation, as demonstrated by the European Rosetta mission [1].

As part of the UK Space Agency National Space Technology Programme, the University of Bristol and Airbus Defence and Space were awarded the project "Robust and Nonlinear Guidance and Control for Landing on Small Bodies", with the aim to investigate the application of robust control techniques for the design and optimisation of D&L* approaches. Although a generic framework was pursued, the project focused on the Martian moon Phobos, which is among the candidates for an interplanetary sample return mission, not only because of the wide scientific interest to solve unknowns about its formation, but also as technological precursors for future manned and unmanned missions to the Martian system [2].

A conceptual separation between guidance and control logics was taken for the project. Within it, previous studies have been dedicated to the assessment [3, 4] and optimisation [5] of different D&L guidance laws. This paper, on the other hand, is entirely devoted to the design of improved D&L strategies via robust control compensation, enabling the adoption of simpler guidance schemes.

In fact, significant room for improvement can be achieved at this level since current state-of-practice in D&L control compensation is conventionally over-simplified or even non-existent, which is only practical if the target body is well-known or if guidance algorithms do not rely on its natural dynamics. The application of robust control techniques is nowadays a well-established practice for spacecraft attitude (rotational) control design [6, 7, 8], but not so much for orbital (translational) control. This paper aims therefore to address the latter point.

The proposed design approach is further motivated by, and based on, the recently-developed structured $\mathcal{H}_\infty$ optimisation paradigm [9, 10], which is particularly well-oriented towards the industrial state-of-practice. This point has been central during the whole project, since the applicability of legacy knowledge is fundamental within the Space industry. Throughout the paper, the suitability of structured $\mathcal{H}_\infty$ for the D&L problem, as well its main challenges, will be further discussed.

The organisation of this paper is as follows: Sec. 2 introduces the Phobos mission benchmark under consideration, Sec. 3 describes the development of models that capture the D&L dynamics and uncertainties, Sec. 4 is then dedicated to the application of the structured $\mathcal{H}_\infty$ framework to synthesise robust control compensators and Sec. 5 to a thorough robustness assessment of these compensators using the structured singular value μ together with high-fidelity Monte-Carlo simulation. The conclusions are given in Sec. 6.

2. PHOBOS MISSION BENCHMARK

Landing on Phobos is particularly challenging because of its reduced mass (it is 8 orders of magnitude smaller than Mars) and proximity to the red planet (mean orbital altitude about 6000 km), which causes the planet's sphere of influence to end just 3.5 km above Phobos' surface. Hence, there is no possibility for Keplerian orbits around Phobos and the third-body perturbation of Mars cannot be neglected.

---

*See Nomenclature at end of article.

Furthermore, due to the irregular shape and mass distribution of Phobos, the gravity of the moon cannot be accurately accounted for by a spherical field, and requiring to be described using a gravity harmonics (GH) model. In this case, using spherical coordinates (r, θ, ϕ) for distance to barycentre, co-latitude and longitude, R for a reference radius and μ_g for the gravitational constant, the gravity potential is given by:

$$U(r, \theta, \phi) = \frac{\mu_g}{R} \sum_{n=0}^{\bar{n}} \left(\frac{R}{r}\right)^{n+1} \sum_{m=0}^{n} C_n^m(\phi) P_n^m(\cos \theta) \tag{1}$$

where:

$$
\begin{aligned}
P_n^m(x) &= (1 - x^2)^{m/2} \frac{\mathrm{d}^m}{\mathrm{d}x^m} P_n(x) \\
P_n(x) &= \frac{1}{2^n n!} \frac{\mathrm{d}^n}{\mathrm{d}x^n} (x^2 - 1)^n
\end{aligned}
\tag{2}
$$

are the associated Legendre polynomials and:

$$C_n^m(\phi) = C_{n,m} \cos m\phi + S_{n,m} \sin m\phi \tag{3}$$

is the expansion of the GH coefficients $C_{n,m}$ and $S_{n,m}$. For Phobos, $\bar{n} = 4$ (i.e., 28 coefficients) is assumed to suffice. However, 19 of these coefficients are highly inaccurately known. In fact, each of them is affected by a Gaussian dispersion with standard deviation equal to its mean value. For an elaborate discussion on this topic, the reader is referred to [3].

Given the small eccentricity of Phobos' orbit around Mars (approximately 0.015), the nonlinear dynamics of a spacecraft in the vicinity of Mars-Phobos system is typically described as a circular restricted three-body problem. In this case, its motion can be written in a body-centred body-fixed (BCBF) frame with origin at the moon's barycentre as:

$$
\begin{bmatrix} \dot{\mathbf{r}}(t) \\ \dot{\mathbf{v}}(t) \\ \dot{\nu}(t) \end{bmatrix} = \mathbf{f}\left(\mathbf{r}(t), \mathbf{v}(t), \nu(t)\right) + \begin{bmatrix} 0_{3\times3} \\ I_{3\times3} \\ 0_{1\times3} \end{bmatrix} \mathbf{a}(t)
\tag{4}
$$

where the state vector $\begin{bmatrix} \mathbf{r}(t) & \mathbf{v}(t) & \nu(t) \end{bmatrix}^{\mathrm{T}}$ gathers the position and velocity of the spacecraft with respect to Phobos, as well as the true anomaly of the latter around Mars, and the control vector $\mathbf{a}(t)$ represents any propulsive acceleration generated by the spacecraft in the BCBF frame. The effects of a complex and uncertain gravity field enter in $\mathbf{f}\left(\mathbf{r}(t), \mathbf{v}(t), \nu(t)\right)$ through the dynamics of $\mathbf{v}$ (where the dependence on t was dropped for clarity), which are given as:

$$\dot{\mathbf{v}} = \boldsymbol{\nabla} U_{\mathrm{SC}}^{\mathrm{Pho}}(\mathbf{r}) + \boldsymbol{\nabla} U_{\mathrm{SC}}^{\mathrm{Mars}}(\mathbf{r}, \nu) - \boldsymbol{\nabla} U_{\mathrm{Pho}}^{\mathrm{Mars}}(\nu) - \omega_{\mathrm{Pho}} \times (\omega_{\mathrm{Pho}} \times \mathbf{r}) - 2\omega_{\mathrm{Pho}} \times \mathbf{v} \tag{5}$$

In this equation, the first three terms represent respectively the gravity of Phobos on the spacecraft (SC), Mars on the spacecraft and Mars on Phobos, computed as a function of the position $\mathbf{r}$ of the spacecraft relative to Phobos and of the latter's location around Mars (via its true anomaly ν). These calculations rely on the potential description of Eq. (1) and become particularly complex due to the frame transformations involved, as detailed in [11]. The remaining terms account for the non-inertial acceleration caused by the fact that the BCBF frame is rotating with Phobos at rate ω_{Pho}.

The generic Space D&L problem is then depicted in Fig. 1 for planar motion but without loss of generality. It describes a spacecraft approaching a moving (small) body target subject to the influence of another (large) body. For this problem, it is assumed that the spacecraft has a dedicated attitude control system that maintains a nadir pointing during D&L (it is a requirement of visual navigation). Coupling effects are then considered at actuator level by reserving a fraction of the available thruster authority (10 to 20%) for attitude control.
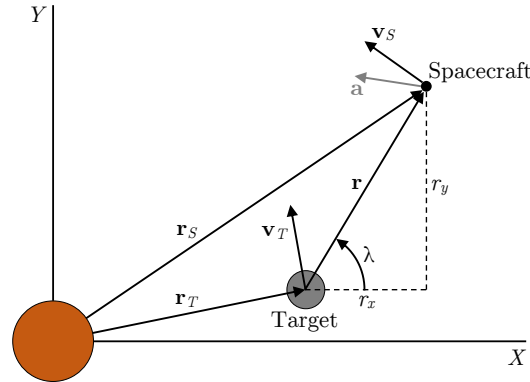

Figure 1. Space D&L problem geometry

Based on this figure, the D&L problem lies on the computation of the acceleration input $\mathbf{a}(t)$ between $t = t_0$ and $t = t_f$ that must be able to:

- Bring the relative position and velocity from the initial boundary conditions $\mathbf{r}(t_0) = \mathbf{r}_0$ and $\mathbf{v}(t_0) = \mathbf{v}_0$ to the final conditions $\mathbf{r}(t_f) = \mathbf{r}_f$ and $\mathbf{v}(t_f) = \mathbf{v}_f$;
- Cope with the effect of uncertainties and external perturbations.

In this problem, the duration from a given instant of time t until the end of the manoeuvre is known as time-to-go, $t_{\mathrm{go}}(t) = t_f - t$, the norm of the relative velocity vector $||\mathbf{v}(t)|| = ||\mathbf{v}_S(t) - \mathbf{v}_T(t)||$ is known as closing speed and the direction from target to spacecraft is known as line-of-sight $\lambda(t)$. For further details on the Space D&L problem, the reader is referred to [4, 12]. Also in these references, the application of closed-loop guidance techniques for the computation of appropriate acceleration laws is addressed.

The focus of the present work, however, is not specifically on closed-loop guidance, but rather on the design of improved D&L strategies via robust control compensation, which ultimately will enable the adoption of (simpler) open-loop guidance schemes. The benchmark architecture is depicted in Fig. 2.
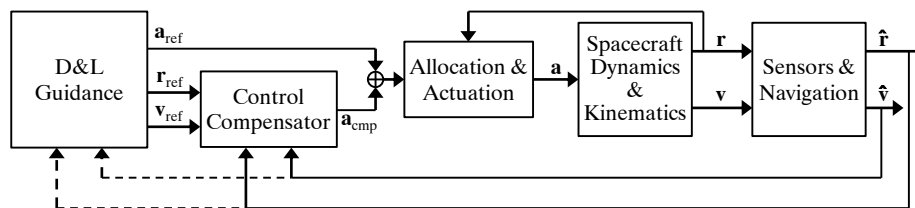

Figure 2. Benchmark architecture (**- - - -** indicates information flow for closed-loop guidance)

In this benchmark, the behaviour of the spacecraft is simulated using three high-fidelity modules, "spacecraft dynamics & kinematics" (SDK), "allocation & actuation" and "sensors & navigation", which have been developed by Airbus within the frame of the project based on their Space operational experience. These modules are part of their industrial testing, simulation & assessment (TSA) facility and will be employed for system development and verification in case of a future Phobos sample return mission.

The SDK model describes the relative motion of Eq. (4) and Fig. 1. Then, the "allocation & actuation" block introduces thruster realisation errors, which may arise from mounting misalignments or limited control authority (as mentioned before, spacecraft thrusters are also responsible for keeping a nadir pointing). It does not account for thruster modulation nor gas dynamics since the natural frequency of these phenomena is one order of magnitude higher than what is used to appropriately model the orbital manoeuvres. Finally, the "sensors & navigation" module implements the algorithms executed to produce relative position and velocity estimates, $\hat{\mathbf{r}}(t)$ and $\hat{\mathbf{v}}(t)$.

In addition, for the "D&L guidance" logic, two paradigms can be defined: [i] open-loop (or implicit), employed when a reference trajectory $\{\mathbf{r}_{\text{ref}}(t), \mathbf{v}_{\text{ref}}(t)\}$ and thruster profile $\mathbf{a}_{\text{ref}}(t)$ are generated before, and remain unchanged during, the descent; or [ii] closed-loop (or explicit), which refers to the case when the thruster profile is computed in real-time to correct the trajectory based on onboard measurements (as shown in Fig. 2 by the dashed lines).

Regardless of the type of guidance that precedes it, the architecture may be augmented with a "control compensator", which introduces an additional acceleration vector command $\mathbf{a}_{\text{cmp}}(t)$ to further alleviate trajectory errors. As anticipated above, the design of such a compensator is the main focus of this paper. For synthesis and analysis purposes, D&L reference trajectories (RTs), defining the set of design points where the controller must be able to operate, are designed through the following process:

1. Analysis of a set of unstable manifolds originated at Libration Point Orbits (LPOs) in the three-body system that intersect Phobos (Fig. 3a);
2. Selection of the manifolds that reach the moon with higher incidence angle and lower closing speed as initial guesses;
3. Definition of a polynomial acceleration profile to bring the final (touchdown) speed to zero;
4. Optimisation of the initial guesses and acceleration profile via nonlinear programming, with the objective of minimum fuel consumption, given as $\Delta V = \int_{t_0}^{t_f} ||\mathbf{a}(\tau)|| \, \mathrm{d}\tau$.

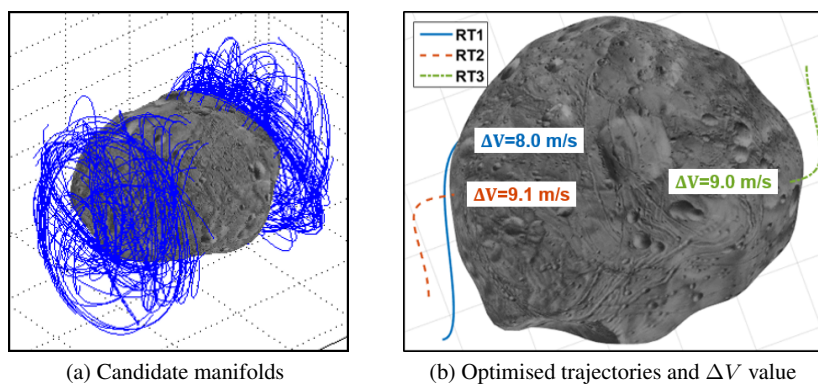(a) Candidate manifolds   (b) Optimised trajectories and ΔV value

Figure 3. Reference D&L trajectory design

To simplify the optimisation runs, a few approximations are considered in the dynamics of the Mars-Phobos system, hence it may not be possible to track the resulting profiles $\{\mathbf{r}_{\text{ref}}(t), \mathbf{v}_{\text{ref}}(t), \mathbf{a}_{\text{ref}}(t)\}$ without any guidance or control feedback loop. For a detailed description on this process, the reader is referred to [3].

For the remainder of this work, the three trajectories visible in Fig. 3b are considered. All of them will be employed for closed-loop validation, but only RT1 will be used for control synthesis purposes.

3. SPACE D&L SUBSYTEM MODELLING

This section introduces the D&L subsystem models developed for control design. It starts with the linearisation of the high-fidelity dynamics of a spacecraft in the vicinity of Mars-Phobos system (Sec. 3.1), followed by the inclusion of gravitational uncertainties through LFT modelling (Sec. 3.2) and concluding with the characterisation of the actuator and navigation models considered (Sec. 3.3).

3.1. Orbital perturbation model

For control design purposes, it is convenient to have a linear representation of Eq. (4). This is achieved via the application of linearised orbital perturbation theory [13]. According to this theory, state and control variables can be defined at different operating points along a given trajectory as the sum of a reference (desired) value and small perturbations (deviations):

$$
\begin{aligned}
\mathbf{r}(t) &= \mathbf{r}_{\text{ref}}(t) + \delta\mathbf{r}(t) \\
\mathbf{v}(t) &= \mathbf{v}_{\text{ref}}(t) + \delta\mathbf{v}(t) \\
\nu(t) &= \nu_{\text{ref}}(t) + \delta\nu(t) \\
\mathbf{a}(t) &= \mathbf{a}_{\text{ref}}(t) + \delta\mathbf{a}(t)
\end{aligned}
\tag{6}
$$

The dynamics of these perturbations is then approximated by the 1^{st} order terms of the Taylor series expansion of $\mathbf{f}\left(\mathbf{r}(t), \mathbf{v}(t), \nu(t)\right)$ around the reference points:

$$
\begin{bmatrix} \delta\dot{\mathbf{r}}(t) \\ \delta\dot{\mathbf{v}}(t) \\ \delta\dot{\nu}(t) \end{bmatrix} = J_{\mathbf{f}}(t) \begin{bmatrix} \delta\mathbf{r}(t) \\ \delta\mathbf{v}(t) \\ \delta\nu(t) \end{bmatrix} + \begin{bmatrix} 0_{3\times3} \\ I_{3\times3} \\ 0_{1\times3} \end{bmatrix} \delta\mathbf{a}(t)
\tag{7}
$$

where the Jacobian matrix is given by:

$$
J_{\mathbf{f}}(t) = \begin{bmatrix} \dfrac{\partial\mathbf{f}}{\partial\mathbf{r}} & \dfrac{\partial\mathbf{f}}{\partial\mathbf{v}} & \dfrac{\partial\mathbf{f}}{\partial\nu} \end{bmatrix} \Bigg|_{\substack{\mathbf{r} = \mathbf{r}_{\text{ref}}(t) \\ \mathbf{v} = \mathbf{v}_{\text{ref}}(t) \\ \nu = \nu_{\text{ref}}(t)}}
\tag{8}
$$

and computed via finite differences due to the complexity of $\mathbf{f}\left(\mathbf{r}(t), \mathbf{v}(t), \nu(t)\right)$. Performing this linearisation at different instants of time t_i, $i = \{1, \ldots, N\}$ along a reference trajectory allows to generate a set of linear time-invariant (LTI) SDK models termed $G_{\text{SDK}}^i(\text{s})$, with the following state-space description:

$$
\left[\frac{\dot{\mathbf{x}}_{\mathrm{SDK}}(\mathrm{s})}{\begin{array}{c}\delta\mathbf{r}(\mathrm{s})\\ \delta\mathbf{v}(\mathrm{s})\end{array}}\right] = \left[\begin{array}{c|c} J_{\mathbf{f}}^i & \begin{array}{c} 0_{3\times3} \\ I_{3\times3} \\ 0_{1\times3} \end{array} \\ \hline I_{6\times6} \quad 0_{6\times1} & 0_{6\times3} \end{array}\right]\left[\frac{\mathbf{x}_{\mathrm{SDK}}(\mathrm{s})}{\delta\mathbf{a}(\mathrm{s})}\right] \tag{9}
$$

where $J_{\mathbf{f}}^i = J_{\mathbf{f}}(t_i)$ and $\mathbf{x}_{\mathrm{SDK}}(\mathrm{s})$ is the internal state vector. The pole-zero map of $G_{\mathrm{SDK}}^i(\mathrm{s})$ along RT1 (refer to Fig. 3b) and the magnitude response of its first channel over time are depicted in Fig. 4.

(a) Pole-zero map   (b) Magnitude response of channel $\delta a_{\mathrm{x}} \to \delta r_{\mathrm{x}}$
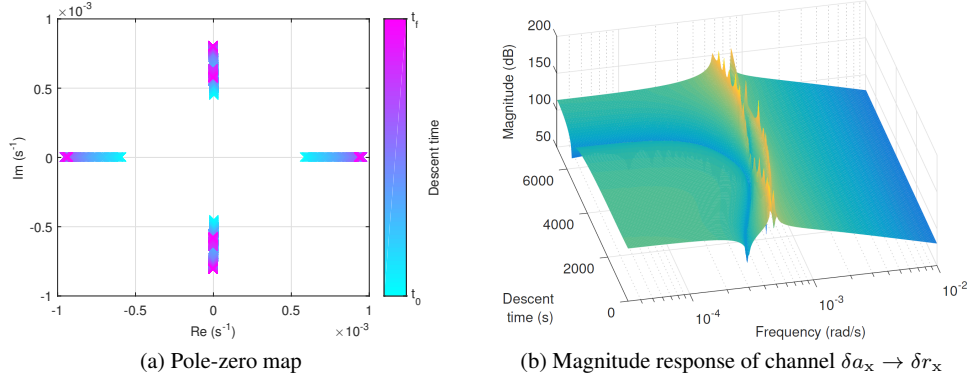
Figure 4. Linear analysis of $G_{\mathrm{SDK}}^i(\mathrm{s})$ for RT1 with $t_i \in [t_0, t_f]$

The shown linear analysis confirms the complexity of the evolution of $G_{\mathrm{SDK}}^i(\mathrm{s})$, with its poles (Fig. 4a) and magnitude peaks (Fig. 4b) changing smoothly but non-monotonically over time and with higher characteristic frequencies towards the final descent time, where the gravity pull of Phobos becomes more intense. It is also confirmed from Fig. 4a that the system is open-loop unstable, since any deviation from the reference trajectory will cause the spacecraft to be pulled towards either Mars or Phobos.

3.2. Inclusion of gravitational uncertainties

As introduced in Sec. 2, 19 out of the 28 GH coefficients are highly inaccurately known. This means that the computation of $J_{\mathbf{f}}^i$ and thus the description of $G_{\mathrm{SDK}}^i(\mathrm{s})$ in Eq. (9) is subject to a high level of uncertainty. To capture the effect of this uncertainty, a mathematical representation known as linear fractional transformation (LFT) is employed.

This type of representation is widely adopted in the robust control framework and thoroughly discussed in [14, 15]. As a cursory example, assume that a given parameter x has a relative uncertainty range w_{x} around its nominal value x_{NOM}, then this can be represented as:

$$
x = x_{\mathrm{NOM}}\left(1 + w_{\mathrm{x}}\delta_{\mathrm{x}}\right), \qquad \delta_{\mathrm{x}} \in [-1, 1] \tag{10}
$$

Following the conventional LFT notation, the product of a signal w with the uncertain parameter δ_{x} can be generically represented as shown in Fig. 5 and written as:

$$
z = xw = \mathcal{F}_{\mathrm{u}}\left\{\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix}, \delta_{\mathrm{x}}\right\}w := \left(M_{22} + M_{21}\delta_{\mathrm{x}}\left(I - M_{11}\delta_{\mathrm{x}}\right)^{-1}M_{12}\right)w \tag{11}
$$

where $\mathcal{F}_u$ is the upper LFT operation, the nominal system $M_{22} = x_{\mathrm{NOM}}$ is retrieved for $\delta_x = 0$ and:

$$M = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ x_{\mathrm{NOM}} w_x & x_{\mathrm{NOM}} \end{bmatrix} \tag{12}$$
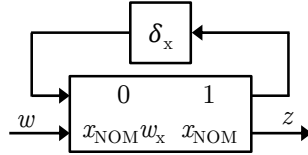

Figure 5. LFT representation of an uncertain parameter

LFTs are particularly attractive due to their extreme modularity and because typical algebraic operations (e.g., inverse, cascade, parallel and feedback connections) preserve the LFT structure. Therefore, in an interconnected system, it is common to isolate what is known as an LTI system, i.e. $M_{22}(s)$, and gather all the "troublemaking" (uncertain, time-varying or nonlinear) components into an operator $\Delta_x = \mathrm{diag}\,(\delta_{x_1}, \delta_{x_2}, \ldots, \delta_{x_n})$ with normalised infinity norm, $||\Delta_x||_\infty \leq 1$. This process is nowadays rather straightforward thanks to MATLAB's Robust Control Toolbox [16].

To build representative LFT models of $G_{\mathrm{SDK}}^i(s)$, a three-step procedure has been adopted. These steps encompass uncertainty selection, plant interpolation and LFT verification and are described in the following paragraphs. The procedure must be repeated for every time instant of interest. Here, 10 design points $i = \{1, \ldots, 10\} : t_i \in [t_0, t_f]$ have been chosen with approximately uniform intervals in terms of closing speed $||\mathbf{v}_{\mathrm{ref}}(t)||$, as evidenced in Fig. 6. The reason for this choice is twofold: [i] information is more refined closer to the final descent time, where the dynamics of $G_{\mathrm{SDK}}^i(s)$ becomes faster (recalling Fig. 4a), and [ii] the evolution of $||\mathbf{v}_{\mathrm{ref}}(t)||$ seems to be determining in the response of the system (comparing Fig. 6 and 4b).
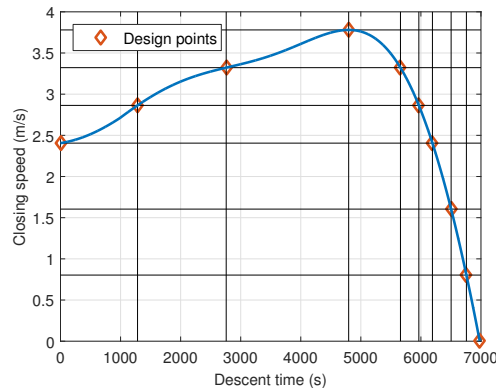

Figure 6. Closing speed of RT1 and design points chosen

3.2.1. Uncertainty selection

In order to minimise the size of the resulting LFT model, only the GH coefficients of Eq. (3) with higher impact on J_f^i are selected and denoted ρ_{GH}. For this choice, two distinct criteria are applied:

[i] impact of each coefficient on the nonlinear simulation of Eq. (4) and [ii] relative weight of each coefficient on dedicated interpolations (see next step). With the two criteria providing consistent results, the following set of 9 coefficients is selected:

$$\rho_{\mathrm{GH}} = \begin{bmatrix} C_{3,0} & C_{3,1} & C_{3,2} & C_{4,0} & C_{4,1} & C_{4,2} & C_{4,3} & S_{3,3} & S_{4,3} \end{bmatrix} \tag{13}$$

This choice will be further revisited at control analysis stage (Sec. 5.3) based on insights from closed-loop sensitivity analysis.

3.2.2. Plant interpolation

Once the set of uncertainties is selected, an appropriate number of dispersed samples is generated and the Jacobian matrix of Eq. (8) is evaluated for each sample. Depending on the dispersion ranges considered, models with different levels of conservativeness can be obtained. As mentioned in Sec. 2, each GH coefficient is assumed to have a Gaussian distribution with standard deviation equal to its mean value. Here, two cases are considered: 2σ and 3σ dispersions, capturing respectively 95.5% and 99.7% of uncertainty.

For each case, a matrix with polynomial dependence on the uncertain parameters $J_{\mathbf{f}}^i(\rho_{\mathrm{GH}})$ is interpolated. This task is accomplished using the orthogonal least-squares approximation of ONERA's APRICOT library [17].

3.2.3. LFT verification

The final step consists in converting the LTI system of Eq. (9) with $J_{\mathbf{f}}^i(\rho_{\mathrm{GH}})$ into an LFT. This is also an embedded capability of the APRICOT library, which adjusts the necessary repetitions of each coefficient in ρ_{GH} to meet a pre-specified approximation error. In the end, two sets of models, $G_{\mathrm{SDK},2\sigma}^i(\mathrm{s})$ and $G_{\mathrm{SDK},3\sigma}^i(\mathrm{s})$, have been obtained for $i = \{1, \ldots, 10\}$, together with corresponding $\Delta_{\mathrm{SDK},2\sigma}^i(\mathrm{s})$ and $\Delta_{\mathrm{SDK},3\sigma}^i(\mathrm{s})$ blocks. Each block contains the uncertainties $\delta c_{n,m}(\mathrm{s})$ and $\delta s_{n,m}(\mathrm{s})$ (with n and m indicating the selected coefficients shown in Eq. (13)), normalised to [-1,1] and repeated as necessary. The sum of these repetitions determines the size of the block and, the larger it is, the more elaborate the impact of the uncertainties is in the dynamics of the system.

It is therefore reasonable to expect that the uncertainty size changes with the level of conservativeness and also along the trajectory. This behaviour is depicted in Fig. 7a. As the figure shows, the size of $\Delta_{\mathrm{SDK},2\sigma}^i(\mathrm{s})$ is never larger than $\Delta_{\mathrm{SDK},3\sigma}^i(\mathrm{s})$ and they both tend to increase throughout time (i.e., towards Phobos), ranging between a value of 3 (when only part of the 9 uncertainties plays a significant role) to 27 (when all the uncertainties are repeated 3 times).

Finally, the accuracy of the LFTs is verified by comparing frequency response samples obtained from the LFT and from equivalently dispersed LTI realisations. An example of such a comparison is shown in Fig. 7b for the singular values of $G_{\mathrm{SDK},3\sigma}^{10}(\mathrm{s})$. In this case, the accuracy of the LFT (uppermost plot) is confirmed by the good matching with LTI responses (bottom plot).

3.3. Actuator and navigation models

As mentioned in Sec. 2, the "allocation and actuation" block of Fig. 2 is responsible for the introduction of thruster realisation errors. These errors are time and trajectory dependent, but they are assumed to be bounded.

(a) SDK LFT and size of $\Delta_{\mathrm{SDK},2\sigma}^{i}(\mathrm{s})$ and $\Delta_{\mathrm{SDK},3\sigma}^{i}(\mathrm{s})$ along RT1

(b) Singular values comparison of $G_{\mathrm{SDK},3\sigma}^{10}(\mathrm{s})$ from LFT (up) and LTI (bottom) samples
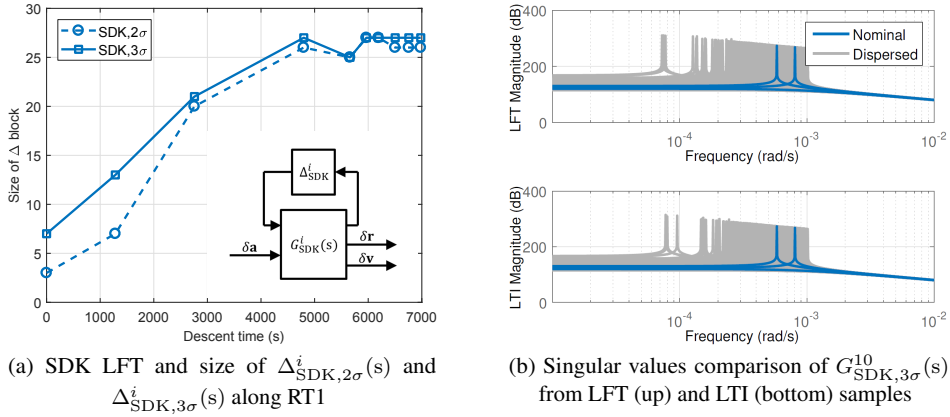
Figure 7. Results of LFT verification step

Actuator errors can then be captured following the same lines of Sec. 3.2, i.e., defining a simple LFT model $G_{\mathrm{A}}(\mathrm{s})$:

$$\mathbf{a}(\mathrm{s}) = \left(I_{3\times3} + \begin{bmatrix} w_{\mathrm{ax}}\delta_{\mathrm{ax}}(\mathrm{s}) & w_{\mathrm{aod}}\delta_{\mathrm{aod}}(\mathrm{s}) & w_{\mathrm{aod}}\delta_{\mathrm{aod}}(\mathrm{s}) \\ w_{\mathrm{aod}}\delta_{\mathrm{aod}}(\mathrm{s}) & w_{\mathrm{ay}}\delta_{\mathrm{ay}}(\mathrm{s}) & w_{\mathrm{aod}}\delta_{\mathrm{aod}}(\mathrm{s}) \\ w_{\mathrm{aod}}\delta_{\mathrm{aod}}(\mathrm{s}) & w_{\mathrm{aod}}\delta_{\mathrm{aod}}(\mathrm{s}) & w_{\mathrm{az}}\delta_{\mathrm{az}}(\mathrm{s}) \end{bmatrix} \right) \mathbf{a}_{\mathrm{cmd}}(\mathrm{s}) \tag{14}$$

and uncertainty block $\Delta_{\mathrm{A}}(\mathrm{s}) = \mathrm{diag}\left(\delta_{\mathrm{ax}}(\mathrm{s}), \delta_{\mathrm{ay}}(\mathrm{s}), \delta_{\mathrm{az}}(\mathrm{s}), I_{3\times3}\delta_{\mathrm{aod}}(\mathrm{s})\right)$. Here, $\delta_{\mathrm{ax}}, \delta_{\mathrm{ay}}, \delta_{\mathrm{az}}$ are uncertainties affecting the three axes and δ_{aod} represents off-diagonal effects. Note that the relative ranges of $\Delta_{\mathrm{A}}(\mathrm{s})$ have been established from Monte-Carlo (MC) nonlinear, time simulations to:

$$w_{\mathrm{ax}} = 0.06, \qquad w_{\mathrm{ay}} = 0.11, \qquad w_{\mathrm{az}} = 0.12, \qquad w_{\mathrm{aod}} = 0.14 \tag{15}$$

and then normalised to [-1,1] with $M(\mathrm{s})$ absorbing these scaling factors.

Since uncertainty ranges are relative, $G_{\mathrm{A}}(\mathrm{s})$ is equally applicable to full acceleration commands $\mathbf{a}_{\mathrm{cmd}}(t)$ and to acceleration command perturbations $\delta\mathbf{a}_{\mathrm{cmd}}(t) = \mathbf{a}_{\mathrm{cmd}}(t) - \mathbf{a}_{\mathrm{ref}}(t)$, as per Eq. (6).

The "sensors & navigation" algorithms, on the other hand, introduce two very different effects: a quantisation and a noise error. The former effect is due to the fact that position and velocity estimates, $\hat{\mathbf{r}}(t)$ and $\hat{\mathbf{v}}(t)$, are updated every 60 s, injecting a non-smooth signal into the system. To attenuate it, these estimates are filtered by a first-order low-pass filter $G_{\mathrm{LPF}}(\mathrm{s})$ with 0.05 Hz bandwidth.

The noise error is accounted for by colouring white noise signals $\mathbf{n}_{\mathrm{r}}(t)$ and $\mathbf{n}_{\mathrm{v}}(t)$ through appropriate transfer functions $G_{\mathrm{rNAV}}(\mathrm{s})$ and $G_{\mathrm{vNAV}}(\mathrm{s})$. These transfer functions have been chosen such that the frequency content produced by the nonlinear "sensors & navigation" block is recovered and correspond to:

$$G_{\mathrm{rNAV}}(\mathrm{s}) = 40\frac{0.001(\mathrm{s}+1)}{\mathrm{s}+0.001}I_{3\times3}, \qquad G_{\mathrm{vNAV}}(\mathrm{s}) = \frac{1}{30}\frac{0.03(\mathrm{s}+1)}{\mathrm{s}+0.03}I_{3\times3} \tag{16}$$

Similarly to the actuator model, these filters are equally applicable to full and perturbed estimates since reference values $\mathbf{r}_{\mathrm{ref}}(t)$ and $\mathbf{v}_{\mathrm{ref}}(t)$ are known without error, so $\hat{\mathbf{r}}(t) - \mathbf{r}(t) = \delta\hat{\mathbf{r}}(t) - \delta\mathbf{r}(t)$ and $\hat{\mathbf{v}}(t) - \mathbf{v}(t) = \delta\hat{\mathbf{v}}(t) - \delta\mathbf{v}(t)$.

4. STRUCTURED ROBUST CONTROL SYNTHESIS

This section is dedicated to the synthesis of robust control compensators using structured $\mathcal{H}_\infty$ optimisation and building up on the models developed in Sec. 3. An introduction of the structured $\mathcal{H}_\infty$ paradigm and its strengths in the face of the Space D&L problem are provided in Sec. 4.1. It is then applied in Sec. 4.2 to design a set of control compensators with distinct properties.

4.1. The structured $\mathcal{H}_\infty$ framework

Although classical techniques are still the state-of-practice for Space control design, they offer limited ability to systematically manage the effects of dynamical perturbations and the specification of multi-channel requirements. Consequently, this approach relies on extensive tuning followed by *a posteriori* verification and validation, which renders the whole process very time and cost consuming, with solutions often found via brute-force search.

The general objective of robust control synthesis is to find a linear controller $K(\mathrm{s})$ such that all the desired requirements of a multi-channel LTI plant $P(\mathrm{s})$ are fulfilled, even in the presence of uncertain dynamics $\Delta(\mathrm{s})$, with $||\Delta(\mathrm{s})||_\infty \leq 1$. The controller uses the output vector $\mathbf{y}$ and generates input vector $\mathbf{u}$ to stabilise and control the plant, which also contains a set of exogenous inputs $\mathbf{w}$ (reference and disturbance signals) and regulated outputs $\mathbf{z}$ (error and performance measurements). This setup is typically represented as the generalised interconnection of Fig. 8.
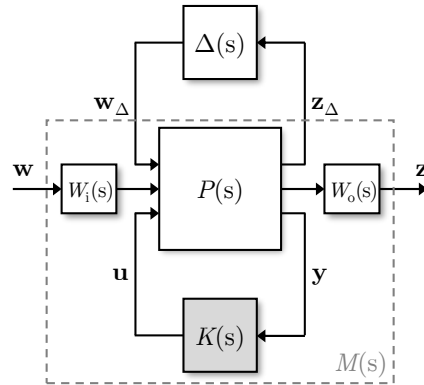

Figure 8. Generalised interconnection in the $\mathcal{H}_\infty$ framework

In this diagram, an input weight $W_\mathrm{i}(\mathrm{s})$ is employed to normalise the frequency content of exogenous signals and an output weight $W_\mathrm{o}(\mathrm{s})$ translates the desired requirements in the frequency domain. Expressing the open-loop dynamics as:

$$\begin{bmatrix} \mathbf{z}_\Delta(\mathrm{s}) \\ W_\mathrm{o}^{-1}(\mathrm{s})\,\mathbf{z}(\mathrm{s}) \\ \hline \mathbf{y}(\mathrm{s}) \end{bmatrix} = \underbrace{\begin{bmatrix} P_{11}(\mathrm{s}) & P_{12}(\mathrm{s}) \\ P_{21}(\mathrm{s}) & P_{22}(\mathrm{s}) \end{bmatrix}}_{P(\mathrm{s})} \begin{bmatrix} \mathbf{w}_\Delta(\mathrm{s}) \\ W_\mathrm{i}(\mathrm{s})\,\mathbf{w}(\mathrm{s}) \\ \hline \mathbf{u}(\mathrm{s}) \end{bmatrix} \tag{17}$$

and closing the interconnection via the lower LFT:

$$\mathcal{F}_l(P, K) := P_{11} + P_{12} K \left(I - P_{22} K \right)^{-1} P_{21} \tag{18}$$

where dependence on s was dropped for clarity, the input-output transfer function $M(s)$ is given by:

$$\begin{bmatrix} \mathbf{z}_\Delta \\ \mathbf{z} \end{bmatrix} = \begin{bmatrix} I & 0 \\ 0 & W_{\text{o}} \end{bmatrix} \mathcal{F}_l(P, K) \begin{bmatrix} I & 0 \\ 0 & W_{\text{i}} \end{bmatrix} \begin{bmatrix} \mathbf{w}_\Delta \\ \mathbf{w} \end{bmatrix} = \underbrace{\begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix}}_{M(\text{s})} \begin{bmatrix} \mathbf{w}_\Delta \\ \mathbf{w} \end{bmatrix} \tag{19}$$

The robust control problem consists then in finding a stabilising controller $K^*(\text{s})$ that minimises the $\mathcal{H}_\infty$-norm of $M(s)$:

$$\min ||M(\text{s})||_\infty = \min \sup_{\omega \in \mathbb{R}} \bar{\sigma}\left(M(j\omega)\right) \tag{20}$$

which corresponds mathematically to its maximum singular value $\bar{\sigma}\left(M(j\omega)\right)$ and physically to the worst-case amplification of exogenous energy-bounded signals.

The output weight $W_{\text{o}}(\text{s})$ normalises the system such that, if all the control requirements are fulfilled:

$$||M(\text{s})||_\infty < 1 \tag{21}$$

and, under the most basic formulation of the problem known as mixed sensitivity [15], $W_{\text{o}}(\text{s})$ is a block-diagonal weight that penalises tracking errors, plant signals and control effort.

This standard $\mathcal{H}_\infty$ synthesis problem is nowadays a well-established practice for Space applications where robustness and performance trade-offs are key [7, 8, 18, 19], but a few practical limitations have slowed down its adoption at an industrial scale. These limitations are mainly related to the inability to constrain $\mathcal{H}_\infty$ controller dimension and structure by design.

The structured $\mathcal{H}_\infty$ paradigm has been proposed as an alternative to tackle the aforementioned limitations in [20]. Its underlying principle is to represent the desired architecture of the controller as an LFT of low-level tunable elements. Since any interconnection of LFTs preserves its structure, the system of Fig. 8 can be rearranged by isolating the tunable elements into a block-diagonal $\tilde{K}(\text{s})$ and absorbing the rest of the structure into $\tilde{P}(\text{s})$. The structured $\mathcal{H}_\infty$ problem consists then in tuning the low-level elements by minimising the global transfer function of Eq. (20). Moreover, in a similar fashion, it is possible to rearrange the plant $\tilde{P}(\text{s})$ and weights $W_{\text{i}}(\text{s})$ and $W_{\text{o}}(\text{s})$ so as to map different combinations of input-output channels and, most notably, individual design requirements. These combinations are then encapsulated as $M(\text{s}) = \text{diag}\left[M_1(\text{s}), \ldots, M_N(\text{s})\right]$ and the synthesis problem becomes:

$$\min \max \left(||M_1(\text{s})||_\infty, \ldots, ||M_N(\text{s})||_\infty\right) \tag{22}$$

enabling the direct handling of multiple requirements, as well as design plants. Furthermore, this algorithm is nowadays able to account for parametric plant uncertainties [10].

Naturally, these advantages come at the expense of a challenging (non-smooth) mathematical problem and, to solve it, structured $\mathcal{H}_\infty$ employs local optimisation methods. To mitigate the local nature of the optimiser, multiple runs are often performed from random starting points. This represents a key issue within the Space industry due to certification concerns and also a breakdown in the design learning experience (i.e., the assessment of system behaviour changes due to changes in the posing of the problem). Hence, the optimiser initialisation and the choice of the parameters that are free to be tuned may become critical for a successful control design. This will be further discussed in Sec. 4.2.

The structured $\mathcal{H}_\infty$ algorithm is part of the *hinfstruct* and *systune* routines of MATLAB [9], with the latter including an easier quantification of $\mathcal{H}_\infty$ requirements and handling of multiple control requirements, channels and models. However, despite the apparent ease of use of software, there are numerous options available to set the tuning requirements and design conditions (sometimes seemingly fighting each other, as characteristic of multiple minima optimisation). This wide availability of choice is a risk in the sense that it provides a confidence level that can complicate the design process by over-constraining the optimisation problem. In this article, this risk is mitigated by posing the problem as in the standard $\mathcal{H}_\infty$ framework and focusing on frequency-domain weighting functions to impose the desired closed-loop properties.

Despite the challenges mentioned above, the effectiveness of structured robust control has been proven through several simulation design problems, e.g., [21, 22], two Space-flown missions [1, 23] and recently in piloted flight tests [24]. The main strength of structured $\mathcal{H}_\infty$ under the scope of these applications (and, indeed, of Space D&L compensation) is its capability to synthesise controllers that are valid not only for one, but for a set of operating points (i.e., plants).

Finally, an additional strength of structured $\mathcal{H}_\infty$ for Space missions lies on its ability to specify controller structure, which allows [i] control gains to be changed *ad hoc* if any re-tuning is required and [ii] industrial legacy knowledge to be kept. An excellent example of the former was demonstrated in [1], where structured $\mathcal{H}_\infty$ was employed for the refinement Rosetta's controller after thruster authority degradation, which was uploaded to the spacecraft before its final insertion manoeuvre with the target comet in May 2014. Taking advantage of legacy knowledge is also a very important point as it enables the transfer of this design approach (with better robust stability and performance properties) to industry without representing a disruptive change for the design teams.

4.2. Application of structured $\mathcal{H}_\infty$ to D&L compensation

As introduced in Sec. 2, the objective of the D&L control compensator is to provide an additional acceleration command $\mathbf{a}_{\mathrm{cmp}}(t)$ to compensate for deviations with respect to a given reference trajectory. In this article, candidate control compensators are designed through a four-step procedure that relies on structured $\mathcal{H}_\infty$ optimisation to exploit the advantages highlighted in Sec. 4.1. This procedure is described in the following paragraphs.

4.2.1. Control structure

To fulfil the control objective mentioned above, the internal structure chosen for the Space D&L control compensator (as introduced in Fig. 2) is evidenced in Fig. 9.
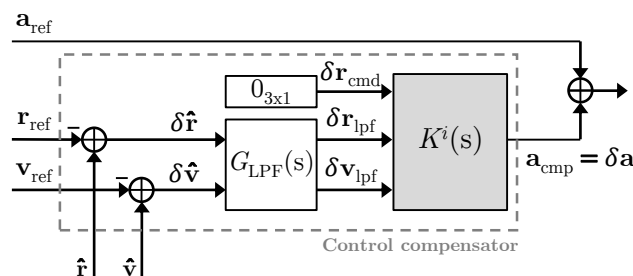

Figure 9. Control compensator architecture

At the core of the compensator, there is an LTI controller $K^i(\mathrm{s})$ that must be designed to track commanded deviations $\delta\mathbf{r}_{\mathrm{cmd}}(t)$ using position and velocity deviation measurements $\delta\hat{\mathbf{r}}(t)$ and $\delta\hat{\mathbf{v}}(t)$ relative to the RT $\{\mathbf{r}_{\mathrm{ref}}(t), \mathbf{v}_{\mathrm{ref}}(t)\}$. These signals are filtered by a first-order low-pass filter $G_{\mathrm{LPF}}(\mathrm{s})$, introduced in Sec. 3.3, to attenuate the quantisation effect injected by the navigation algorithms.

An additional velocity command $\delta\mathbf{v}_{\mathrm{cmd}}(t)$ could have been included, but it was found to be redundant and hence not considered. For this case, the commanded deviation signal is fixed to zero, $\delta\mathbf{r}_{\mathrm{cmd}}(t) = 0$, but it can also be employed to command small changes of landing site without the need for redesigning the trajectory.

With such a structure (fed by position and velocity deviations and providing an acceleration command $\mathbf{a}_{\mathrm{cmp}}(t)$ that is summed to a reference profile $\mathbf{a}_{\mathrm{ref}}(t)$) the controller can be designed based on the orbital perturbation model of Sec. 3.1 by realising that the acceleration command is actually a perturbation, i.e., $\mathbf{a}_{\mathrm{cmp}}(t) = \delta\mathbf{a}(t)$. The closed-loop interconnection is then depicted in Fig. 10.
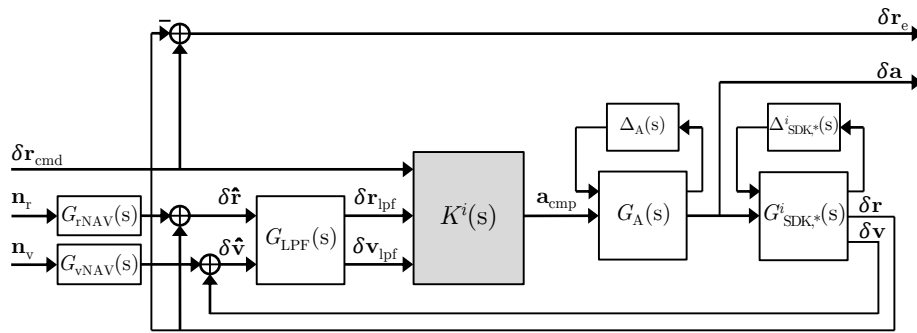

Figure 10. Closed-loop model for control synthesis

In addition to the blocks already present in the previous figure, Fig. 10 features the actuator model $G_{\mathrm{A}}(\mathrm{s})$, uncertainty $\Delta_{\mathrm{A}}(\mathrm{s})$ and navigation noise-shaping filters $G_{\mathrm{rNAV}}(\mathrm{s})$ and $G_{\mathrm{vNAV}}(\mathrm{s})$, all of them described in Sec. 3.3, as well as the orbital perturbation models $G_{\mathrm{SDK},*}^i(\mathrm{s})$ and gravitational uncertainty $\Delta_{\mathrm{SDK},*}^i(\mathrm{s})$ developed in Sec. 3.2. Here, the superscript $i = \{1, \ldots, 10\}$ specifies different designs point along the RT and $*$ represents 2σ or 3σ, depending on the LFT conservativeness at each point.

Generically speaking, a different LTI controller $K^i(\mathrm{s})$ can be synthesised and implemented for each design point i. Nevertheless, *hinfstruct* offers additional capabilities including multi-plant design (when a single controller is synthesised so as to fulfil the control requirements at all the design points simultaneously) and self-scheduled design (similar to multi-plant, but with the controller parameterised as a function of a scheduling variable). A thorough review of these alternatives is provided in [21, 22]. The multi-plant controller synthesis approach proved to be very successful for this D&L problem. Hence, in Fig. 9 and 10 there will be only one controller $K^i(\mathrm{s}) = K(\mathrm{s})$ for all the design points $i = \{1, \ldots, 10\}$.

4.2.2. Control requirements

In order to ensure a successful control compensation, the controller $K(\mathrm{s})$ must be designed in such a way that, even under an uncertain gravitational environment and actuator responses, two driving requirements are fulfilled: [i] *tracking*, minimising the error $\delta\mathbf{r}_{\mathrm{e}}(\mathrm{s}) = \delta\mathbf{r}_{\mathrm{cmd}}(\mathrm{s}) - \delta\mathbf{r}(\mathrm{s})$ and [ii] *actuation*, minimising the overall control effort $\delta\mathbf{a}(\mathrm{s})$.

As mentioned before, control requirements are posed by rearranging the design problem into the generalised framework of Fig. 8. More specifically, comparing the latter figure with Fig. 10, one can identify as exogenous inputs $\mathbf{w}(\mathrm{s}) = [\delta\mathbf{r}_{\mathrm{cmd}}(\mathrm{s}); \mathbf{n}_{\mathrm{r}}(\mathrm{s}); \mathbf{n}_{\mathrm{v}}(\mathrm{s})]$ and as regulated outputs $\mathbf{z}(\mathrm{s}) = [\delta\mathbf{r}_{\mathrm{e}}(\mathrm{s}); \delta\mathbf{a}(\mathrm{s})]$. In addition, all the uncertain elements are combined into a single uncertainty block $\Delta^i(\mathrm{s}) = \mathrm{diag}\left(\Delta^i_{\mathrm{SDK},*}(\mathrm{s}), \Delta_{\mathrm{A}}(\mathrm{s})\right)$ and everything else is gathered in the LTI system $M^i(\mathrm{s})$, again for $i = \{1, \ldots, 10\}$.

Control requirements are then defined through the appropriate choice of frequency-dependent weights $W_{\mathrm{i}}(\mathrm{s})$ and $W_{\mathrm{o}}(\mathrm{s})$, also contained in $M^i(\mathrm{s})$ as shown in Fig. 8. In fact, weight selection is a key factor for the synthesis. For standard $\mathcal{H}_\infty$ synthesis, weight complexity much be chosen as low as possible to limit the controller order, but this does not apply with structured $\mathcal{H}_\infty$. In addition, thanks to the multi-plant approach of the latter, it is possible to use different weights for different design points i. Still, the number of weights shall be the lowest possible in order to reduce the designer's work (i.e., fewer parameters to tune) and the risk of over-constraining the optimisation with conflicting requirements. Hence, for this problem, weights were chosen to be the same for all the design points (in general, this might not be possible).

The input weight $W_{\mathrm{i}}(\mathrm{s})$ is employed for a differential scaling within the input signal $\mathbf{w}(\mathrm{s})$. For this specific case, with main focus on the effect of gravitational uncertainties, no noise has been considered for design, therefore $W_{\mathrm{i}}(\mathrm{s}) = \mathrm{diag}\left(I_{3\times3}, 0_{3\times3}, 0_{3\times3}\right)$. Resilience to noise is achieved by adjusting these values accordingly. The output weight $W_{\mathrm{o}}(\mathrm{s})$ is applied to $\mathbf{z}(\mathrm{s})$ and partitioned as $W_{\mathrm{o}}(\mathrm{s}) = \mathrm{diag}\left(W_S(\mathrm{s}), W_A(\mathrm{s})\right)$, where $W_S(\mathrm{s})$ imposes *tracking* requirements via $\delta\mathbf{r}_{\mathrm{e}}(\mathrm{s})$ and $W_A(\mathrm{s})$ *actuation* requirements through $\delta\mathbf{a}(\mathrm{s})$.

More specifically, as a consequence of Eq. (21), the magnitude response of $\delta\mathbf{r}_{\mathrm{e}}(\mathrm{s})$ for a unitary command in $\delta\mathbf{r}_{\mathrm{cmd}}(\mathrm{s})$ is bounded by $W_S^{-1}(\mathrm{s})$. This inverse shall then have [i] small low-frequency gain (10^{-3}) for little steady-state error, [ii] reasonable high-frequency gain (2) for good stability margins, [iii] a roll-over frequency (10^{-2} rad/s) which is appropriate for the dynamics of the problem and [iv] zero off-diagonal terms to minimise cross-coupled interactions. Hence, $W_S(\mathrm{s})$ can be defined as:

$$W_S(\mathrm{s}) = \frac{\frac{1}{2}\mathrm{s} + 10^{-2}}{\mathrm{s} + 10^{-5}} \, I_{3\times3} \tag{23}$$

In a similar fashion, the magnitude response of $\delta\mathbf{a}(\mathrm{s})$ for a unitary command in $\delta\mathbf{r}_{\mathrm{cmd}}(\mathrm{s})$ is bounded by $W_A^{-1}(\mathrm{s})$. This inverse shall then have [i] reasonable low-frequency gain (5.5×10^{-4}) which establishes the maximum control effort, [ii] small high-frequency gain (10^{-9}) for little reactivity to noisy signals, [iii] a roll-off frequency (10^{-3} rad/s) able to accommodate the *tracking* bandwidth and again [iv] zero off-diagonal terms to minimise cross-coupled interactions. Therefore:

$$W_A(\mathrm{s}) = \frac{\mathrm{s} + \frac{10}{5.5}}{10^{-9}\,\mathrm{s} + 10^{-3}} \, I_{3\times3} \tag{24}$$

4.2.3. Control synthesis

Following the multi-plant approach, plants $M^i(\mathrm{s})$ are aggregated in a block-diagonal structure and the *hinfstruct* routine is called to find a single stabilising controller $K(\mathrm{s})$ that meets the requirements at every LTI point $i = \{1, \ldots, 10\}$.

As mentioned in Sec. 4.1, *hinfstruct* is nowadays able to account for parametric uncertainties $\Delta^i(s)$ in the optimisation problem. In the same section, it was anticipated that the initialisation and choice of the control parameters that are free to be tuned are also key factors for the success of the optimisation. Without a proper initialisation, i.e., using random initial guesses, the optimiser is often unable to find a solution or ends up converging towards a controller with fast and/or unstable poles, which is not practical to implement. Hence, assuming a third-order controller and partitioning it as:

$$
\left[\begin{array}{c} \dot{\mathbf{x}}_{\mathrm{K}}(s) \\ \hline \mathbf{a}_{\mathrm{cmp}}(s) \end{array}\right] = \left[\begin{array}{c|ccc} A_{\mathrm{K}} & B1_{\mathrm{K}} & B2_{\mathrm{K}} & B3_{\mathrm{K}} \\ \hline C_{\mathrm{K}} & D1_{\mathrm{K}} & D2_{\mathrm{K}} & D3_{\mathrm{K}} \end{array}\right] \left[\begin{array}{c} \mathbf{x}_{\mathrm{K}}(s) \\ \hline \delta\mathbf{r}_{\mathrm{cmd}}(s) \\ \delta\mathbf{r}_{\mathrm{lpf}}(s) \\ \delta\mathbf{v}_{\mathrm{lpf}}(s) \end{array}\right] \tag{25}
$$

initial guesses were chosen as follows:

- $A_{\mathrm{K}} = -0.01\,I_{3\times3}$ for stable poles with reasonable frequency;
- $B1_{\mathrm{K}} = -B2_{\mathrm{K}} = I_{3\times3}$ and $B3_{\mathrm{K}} = 0_{3\times3}$ to lead the optimiser towards the consideration of the error $\delta\mathbf{r}_{\mathrm{cmd}}(s) - \delta\mathbf{r}_{\mathrm{lpf}}(s)$;
- $C_{\mathrm{K}} = D2_{\mathrm{K}} = D3_{\mathrm{K}} = I_{3\times3}$, respectively responsible for integrative, proportional and derivative control actions;
- $D1_{\mathrm{K}} = 0_{3\times3}$, which should not be needed since control compensation is independent of the reference trajectory.

This choice results in 72 parameters to be tuned. However, if all of them are left free, the optimiser tends to find physically inappropriate solutions, featuring very different dynamics on each control axis or strong cross-coupled interactions between them. This was solved by fixing to zero all the off-diagonal terms of the sub-matrices in Eq. (25), reducing the number of parameters to 24. As a consequence, the highly-coupled D&L dynamics will be tackled by a diagonal (but robust) controller.

With this strategy, four controllers have been designed, each of them accounting for a different set of uncertainties (recall Fig. 10). This selection is summarised in Table I. For better results, and referring to the latter table, the controller $K_{\mathrm{A}}(s)$ has been employed as initial condition for the synthesis of $K_{\mathrm{A},2\sigma}(s)$ and $K_{\mathrm{A},3\sigma}(s)$.

Table I. Specification of the designed controllers

Controller	Uncertainties considered
$K_0(s)$	None (non-robust controller)
$K_{\mathrm{A}}(s)$	$\Delta_{\mathrm{A}}(s)$ only
$K_{\mathrm{A},2\sigma}(s)$	$\Delta_{\mathrm{A}}(s)$ and $\Delta^i_{\mathrm{SDK},2\sigma}(s)$
$K_{\mathrm{A},3\sigma}(s)$	$\Delta_{\mathrm{A}}(s)$ and $\Delta^i_{\mathrm{SDK},3\sigma}(s)$

4.2.4. Nominal results

Before proceeding to an in-depth robustness analysis of the designed controllers, preliminary tests are conducted with models in nominal conditions. First, an LTI analysis and comparison of closed-loop frequency responses using the four controllers of Table I is provided in Fig. 11. For each case, the figure shows the singular values for the $\delta\mathbf{r}_{\mathrm{e}}(s)$ and $\delta\mathbf{a}(s)$ transfer functions (uppermost and bottom plots, respectively), the corresponding $\mathcal{H}_\infty$ constraints $W_S^{-1}(s)$ and $W_A^{-1}(s)$, and the

response evolution along RT1 (using different colours, similarly to Fig. 4a). In addition, a distinction is made between diagonal and off-diagonal terms (continuous and dashed lines), which shall be as separated as possible to minimise cross-coupled interactions.
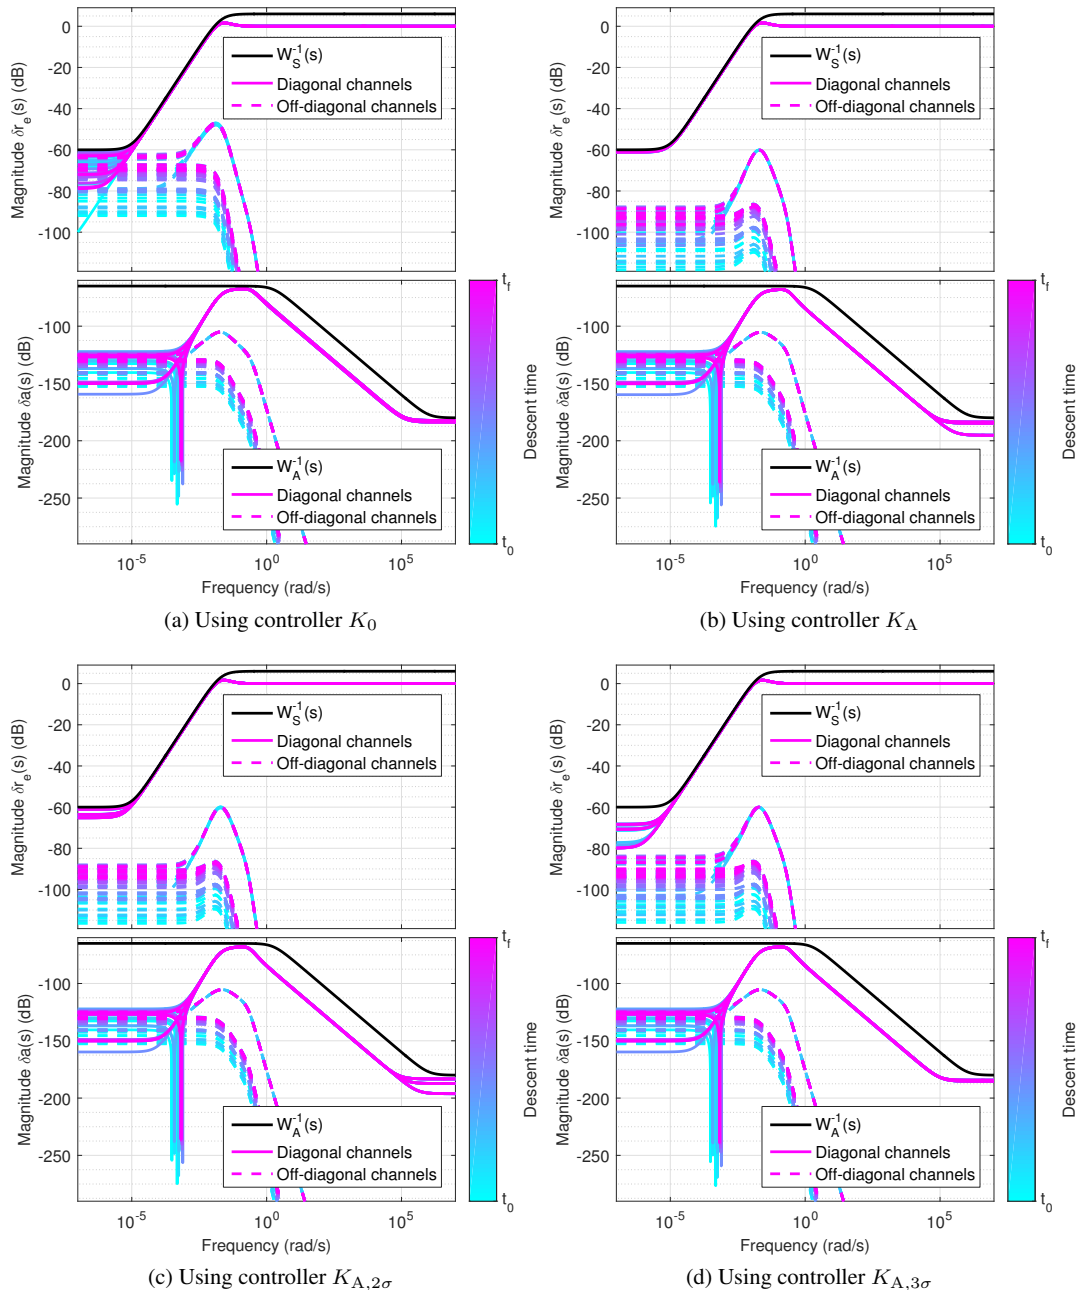

Figure 11. Multi-channel singular values along RT1 and $t_i \in [t_0, t_f]$ in nominal conditions ($\Delta(\mathrm{s}) = 0$) against $W_S^{-1}(\mathrm{s})$ (uppermost plot) and $W_A^{-1}(\mathrm{s})$ (bottom plot)

From this figure, it is clear that, in nominal conditions, the four controllers meet the design requirements for all the points throughout RT1. It is also visible that margins with respect to $W_S^{-1}(\mathrm{s})$ and $W_A^{-1}(\mathrm{s})$ become tighter towards the final descent time, especially for off-diagonal channels. In

fact, for the non-robust controller K_0 (Fig. 11a), there is not enough separation between diagonal and off-diagonal terms, which anticipates robustness issues. This separation is then achieved with the other controllers, which account for cross-couplings introduced by the actuators and/or by gravitational uncertainties. These controllers (Fig. 11b to 11d) show very similar responses (indeed, major differences are only visible using robust performance analysis or highly dispersed simulations), although it is observed that they have sequentially smaller diagonal steady-state errors, at the expense of larger off-diagonal components.

In addition, to demonstrate the effect of control compensation, a nonlinear simulation of RT1 is depicted in Fig. 12, showing reference and actual trajectories, $\mathbf{r}_{\mathrm{ref}}(t)$ and $\mathbf{r}(t)$, as well as arrows with magnitude and direction of acceleration reference $\mathbf{a}_{\mathrm{ref}}(t)$, compensation $\mathbf{a}_{\mathrm{cmp}}(t)$, and total signal (i.e., $\mathbf{a}_{\mathrm{ref}}(t) + \mathbf{a}_{\mathrm{cmp}}(t)$) at different points.

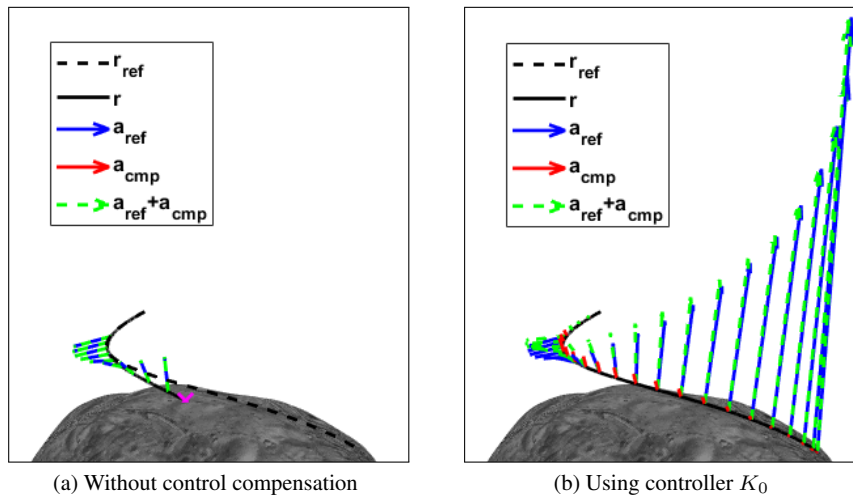(a) Without control compensation (b) Using controller K_0

Figure 12. Impact of control compensation for RT1 in nominal conditions

On the left-hand side, Fig. 12a, no control compensation is used, i.e., $\mathbf{a}_{\mathrm{cmp}}(t) = 0$, so RT1 cannot be followed and the spacecraft crashes on Phobos before reaching the desired target. This potential failure, i.e., in the absence of any guidance or control feedback loop, had already been identified in Sec. 2. A successful D&L is then executed when the compensator K_0 is introduced, as displayed in Fig. 12b (the same result is obtained with any of the other controllers since they have similar nominal performances). Here, the compensator command $\mathbf{a}_{\mathrm{cmp}}(t)$ is clearly visible, with smaller magnitude than the reference vector, but enough to correct it and track the desired trajectory.

5. ROBUSTNESS ANALYSIS AGAINST UNCERTAINTIES

In this section, the impact of gravitational uncertainties when using different compensators is assessed. This assessment is carried out first via the structured singular value μ (introduced in Sec. 5.1 and applied in Sec. 5.2) and then through nonlinear Monte-Carlo simulation (Sec. 5.4). In addition, closed-loop sensitivity information provided by μ analysis is addressed in Sec. 5.3 as a way to validate and complement the uncertainty selection choice made in Sec. 3.2.

5.1. The μ analysis framework

The fundamental approach for analytical system robustness assessment is based on the structured singular value μ. This operator was introduced in [25] and relies on the LFT representation in the form of Fig. 8 to determine if a system $M(\mathrm{s})$, partitioned as in Eq. (11), is stable in the presence of all the normalised LTI uncertainties $||\Delta(\mathrm{s})||_\infty \leq 1$. This is ensured by the existence of the inverse $(I - M_{11}\Delta)^{-1}$ in Eq. (11), and verified by computing:

$$\mu(M_{11}) = \frac{1}{\min_\Delta \left\{ \bar{\sigma}(\Delta) : ||\Delta||_\infty \leq 1, \det\left(I - M_{11}\Delta\right) = 0 \right\}} \tag{26}$$

with $\mu(M_{11}) \in \mathbb{R}^+$ if a solution exists. The system is then said to have robust stability (RS) if and only if its nominal component M_{22} is stable and [14, 15]:

$$\mu(M_{11}(j\omega)) < 1, \qquad \forall\, \omega \in \mathbb{R} \tag{27}$$

Moreover, the norm of the smallest set of uncertainties that destabilises the system is given by $||\mu(M_{11})||_\infty^{-1}$. The computation of Eq. (26) is known to be non-deterministic polynomial-time hard, hence its estimation relies on lower and upper bounds. In practice, they provide an analytical guarantee of existence of at least one perturbation that equals the lower bound degradation and of inexistence of any perturbation that exceeds the upper bound degradation. For more accurate bounds, the size of LFT models shall be kept as small as possible while capturing the most relevant physical phenomena of the real system and their interplay with the uncertainties [14].

Additional valuable insights can also be extracted using μ analysis [26, 27] including: critical uncertainty combinations throughout frequency (which allow to identify worst-case configurations without the need for random sampling methods such as Monte-Carlo) and sensitivities that quantify the impact of each parameter in the solution of μ. All the necessary algorithms are also included in MATLAB's Robust Control Toolbox [16], although there are other alternatives available, e.g. [28].

Furthermore, the structured singular value μ can also be employed for robust performance analysis, i.e., to check if closed-loop requirements defined as in Eq. (21) remain fulfilled under the effect of all the allowable uncertainties $||\Delta(\mathrm{s})||_\infty \leq 1$. Without uncertainties, nominal performance (NP) is assessed by verifying if Eq. (21) holds for $\Delta(\mathrm{s}) = 0$ over frequency, which is equivalent to:

$$\bar{\sigma}(M_{22}(j\omega)) < 1, \qquad \forall\, \omega \in \mathbb{R} \tag{28}$$

Under the presence of uncertainties, the channel $\mathbf{z} \rightarrow \mathbf{w}$ of Fig. 8 is closed through a fictitious complex uncertainty $\Delta_\mathrm{p}(\mathrm{s})$ and the μ test is now applied to $M(\mathrm{s})$ instead of $M_{11}(\mathrm{s})$, considering the augmented perturbation $\mathrm{diag}\left(\Delta(\mathrm{s}), \Delta_\mathrm{p}(\mathrm{s})\right)$. Similarly to the stability case, the system with requirements defined as in Eq. (21) is said to have robust performance (RP) if and only if it has nominal stability and:

$$\mu(M(j\omega)) < 1, \qquad \forall\, \omega \in \mathbb{R} \tag{29}$$

with smaller values of $\mu(M(j\omega))$ indicating better RP properties.

5.2. Application of μ analysis to D&L compensation

The analysis framework introduced in the previous section can now be applied to assess and compare robust stability and performance properties of D&L compensation systems using the multi-plant controllers synthesised in Sec. 4.2.3 and summarised in Table I. Similarly to the analysis carried out in Sec. 4.2.4, this is achieved by closing the loop with the different multi-plant controllers at every LTI point. Also, the uncertainties considered for analysis do not necessarily need to coincide with the ones accounted for the design of each controller. An overview of all the information provided by μ analysis is illustrated in Fig. 13 for the LTI point $i = 8$, which corresponds to a descent time around 6500 s, for three pairs formed by examining K_A and $K_{A,3\sigma}$ with the uncertain LFTs on Δ_A and $\mathrm{diag}\left(\Delta_A, \Delta^8_{\mathrm{SDK},3\sigma}\right)$. On the uppermost subplots, curves related to the RS ($-\triangledown-$), NP ($-\square-$) and RP ($-\triangle-$) tests of Eq. (27), (28) and (29), are depicted over frequency. For RS and RP, there is actually a dashed and a continuous line representing the lower and upper bound for the computation of μ which, for practical effects, show a very good agreement in this scenario. The bottom subplots then show the sensitivity of RP results against each element of Δ, which is quantified by $\partial\mu(M^i(j\omega))/\partial\delta_x$ and directly computed by the μ algorithms. A higher sensitivity indicates that the corresponding uncertainty has a stronger influence in the level of RP obtained.

The first plot, Fig. 13a, shows the closed-loop analysis with K_A, designed for Δ_A only, against the effect of Δ_A. Since all the curves lie below 1, this implies that NP, RS and RP conditions are met, which is natural as the uncertainties considered for design and analysis are the same. In nominal conditions, NP ($-\square-$) is more demanding at low and very high frequencies, which could also be anticipated from the proximity of the closed-loop responses in Fig. 11 with the W_S^{-1} and W_A^{-1} requirements. Stability degradation, i.e., RS ($-\triangledown-$), is only affected at low frequencies and there is always a significant margin until instability, i.e., until $\mu(M^i_{11}) = 1$.

Finally, RP ($-\triangle-$) envelops, and is driven by, the shapes of NP and RS, which is a clear result of Eq. (11): in the presence of uncertainties, closed-loop performance is constrained by nominal performance (i.e., M^i_{22}) and by its interplay with the uncertainty channels (through M^i_{11}, M^i_{12} and M^i_{21}). In terms of sensitivities (bottom subplot), it is clear that all the uncertain elements have similar trends, but the off-diagonal uncertainty δ_{aod} is the most decisive.

Figure 13b shows the same analysis but with K_A now against the combined effect of actuator and GH uncertainties, i.e. $\mathrm{diag}\left(\Delta_A, \Delta^8_{\mathrm{SDK},3\sigma}\right)$. As expected, considering a wider range of uncertainties leads to the same NP, but degraded robustness properties, particularly in terms of RP. As evidenced by the uppermost result, the RP conditions are no longer met by K_A. Also, the elements of $\Delta^8_{\mathrm{SDK},3\sigma}$ are now included in the sensitivities (bottom) subplot with continuous lines. In fact, the impact of these elements becomes comparable to that of Δ_A, precisely at low frequencies where $\mu(M^8) \approx 1$.

Finally, Fig. 13c shows the results against the same previous set of uncertainties $\mathrm{diag}\left(\Delta_A, \Delta^8_{\mathrm{SDK},3\sigma}\right)$, but now replacing the controller with $K_{A,3\sigma}$, designed for the latter set. This leads to all the NP, RS and RP conditions being met by a change of closed-loop behaviour at low-frequency. Nevertheless, the RP sensitivity frequency responses remain roughly the same, which indicates that this type of insight is not intrinsically dependent on the controller being analysed.

In order to conclude the analysis, RP levels against the complete set of uncertainties have been computed along RT1 using the control compensators not addressed in Fig. 13. Results with K_0 and $K_{A,2\sigma}$ are provided in Fig. 14.
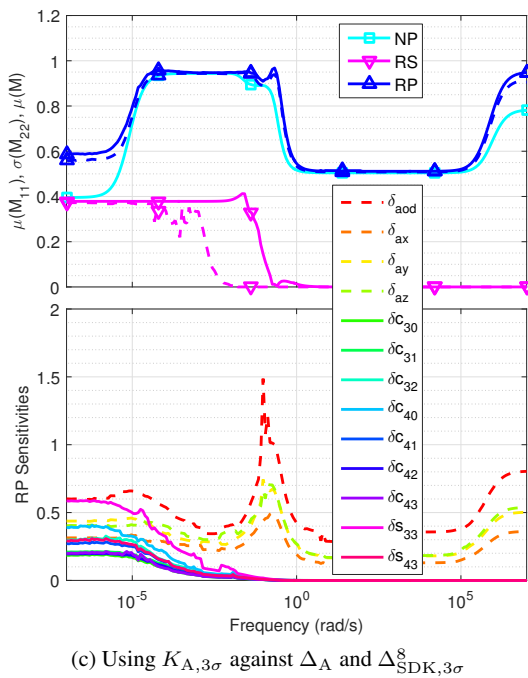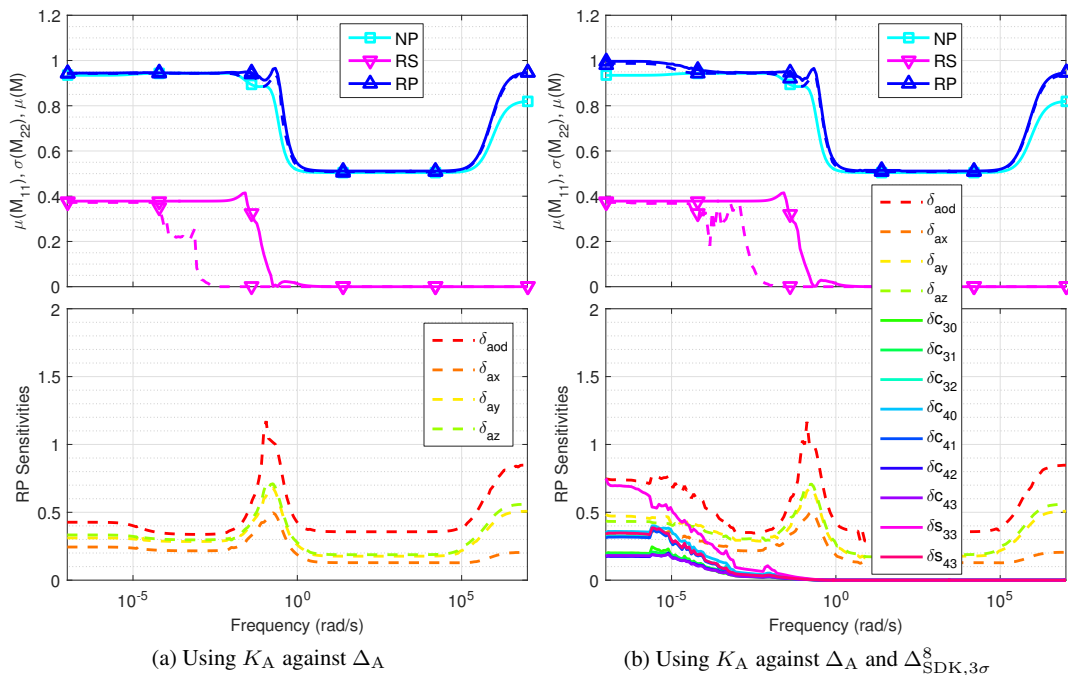
(a) Using K_A against Δ_A

(b) Using K_A against Δ_A and $\Delta_{\mathrm{SDK},3\sigma}^8$

(c) Using $K_{A,3\sigma}$ against Δ_A and $\Delta_{\mathrm{SDK},3\sigma}^8$

Figure 13. NP, RS, RP results (top subplots) and μ sensitivities (bottom subplots) at $i = 8$

It is confirmed from Fig. 14a that K_0, being a non-robust controller, exhibits RP peaks well above 1, hence it is not suitable for the D&L scenario. This is no longer the case with $K_{A,2\sigma}$, in Fig. 14b, whose RP levels lie in-between those of K_A and $K_{A,3\sigma}$. For both cases, as anticipated with the nominal analysis of Fig. 11, performance degradation will become more intense closer to touchdown.
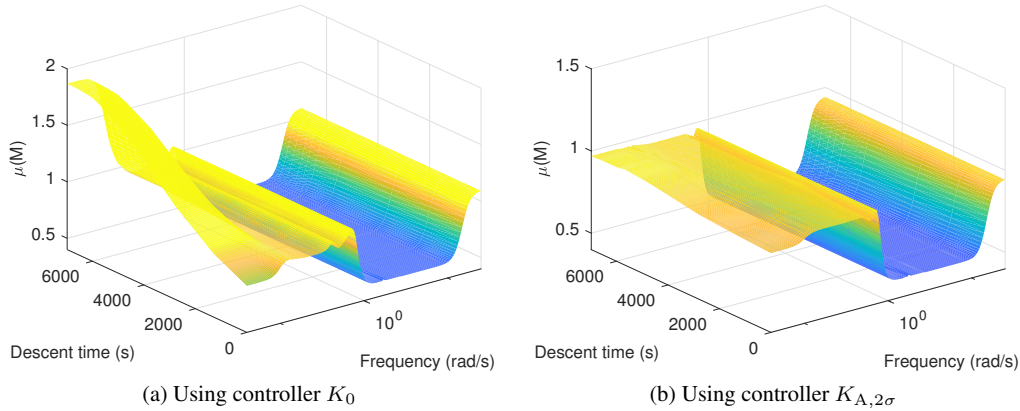
(a) Using controller K_0 (b) Using controller $K_{A,2\sigma}$

Figure 14. RP results against Δ_A and $\Delta_{SDK,3\sigma}^i$ along RT1 with $t_i \in [t_0, t_f]$

5.3. Sensitivity analysis

It was mentioned in Sec. 3.2.1 that an *open-loop* sensitivity analysis was performed to identify the uncertain GH coefficients with higher impact on the dynamics of the problem. More specifically, this consisted of repeating the interpolation of the Jacobian matrices J_f^i for each coefficient separately and counting the number of J_f^i terms where that coefficient accounts for a variation above 10% its nominal value. This allowed to quantify the impact of each parameter and select the 9 most relevant for LFT generation.

Here, this quantification method is revisited using the *closed-loop* sensitivity information provided by μ analysis in Sec. 5.2 and the fact that this type of insight is not intrinsically dependent on the controller used to close the loop (when the controllers are designed to achieve similar objectives).

An overview of the results that can be obtained using the aforementioned considerations is depicted in Fig. 15, for which $K_{A,3\sigma}$ and 3σ uncertainty dispersions were assumed. It shows, per normalised SDK uncertainty δ_x (where x identifies the coefficients selected in Eq. (13)) and as a function of the design points $i = \{1, \ldots, 10\}$, the following: [i] in yellow, the sensitivity peak values computed via μ analysis (e.g., for $i = 8$, these values correspond to the peaks of the bottom plots of Fig. 13c) and [ii] in dark blue, the number of J_f^i terms where δ_x has a relative LFT interpolation weight larger than 10%, normalised so as to have the same scale as the μ sensitivities.

Figure 15 then provides a visual confirmation that the sensitivity insights obtained using two very different methods (open-loop LFT interpolation and closed-loop μ analysis) are quite consistent for this scenario. This is particularly true when identifying cases with smaller sensitivities. But most importantly, the latter figure allows to quantify the relative impact of each GH coefficient on the dynamics of the problem, as well as its evolution through the descent. For example, it is clear that $\delta s_{3,3}$ and $\delta c_{4,0}$ are the most determining in this case, with a growing effect throughout the trajectory, while others, such as $\delta c_{3,2}$ or $\delta c_{4,1}$ are more intense at earlier stages of the descent. LFT models featuring uncertainty sets other than Eq. (13) could also be included in the comparison.

Besides being useful for LFT reduction (as in this article), the determination of the driving GH coefficients has important practical implications. Most notably, it can be employed to support key mission choices (such as trajectory selection) or refined investigation aimed at reducing the uncertainty range of driving coefficients.
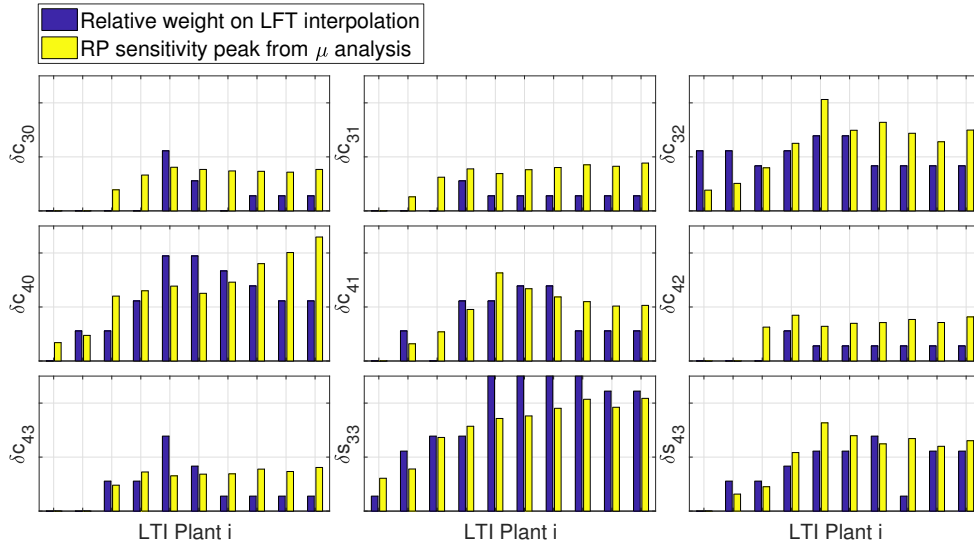
Figure 15. Sensitivity indicators from 3σ LFT interpolation and μ analysis along RT1

5.4. Monte-Carlo verification

The robustness of the four controllers of Table I is now verified using the benchmark introduced in Sec. 2 for the nonlinear high-fidelity simulation of RT1. Each control compensator is tested against the same 2000 Monte-Carlo (MC) samples of the 19 GH coefficients in Eq. (3), randomly sampled with Gaussian distributions. Since the focus of the analysis is on the impact of these coefficients, no additional perturbations from actuators or navigation are included in the simulations.

In addition, special worst-cases (WCs) are also tested for each compensator. These WCs correspond to the combination of uncertainties associated with the RP peaks $\Delta_{\mathrm{SDK},3\sigma}^{i,\mathrm{WC}}$, which is determined via μ analysis at every LTI point $i = \{1,\ldots,10\}$. Since only 9 out of the 19 GH coefficients are captured by the LFTs employed for μ analysis, every combination $\Delta_{\mathrm{SDK},3\sigma}^{i,\mathrm{WC}}$ is tested with the remaining 10 coefficients simultaneously set to their nominal, $+3\sigma$ and -3σ values. Discarding all the repeated combinations, this originates between 18 and 24 WCs per controller.

The outcome of the verification campaigns is depicted in Fig. 16, showing nominal simulations with a dark continuous line, MC runs with lighter lines and the aforementioned WCs with dashed lines. For each control compensator, distance and speed errors with respect to the RT, $|\mathbf{r}(t) - \mathbf{r}_{\mathrm{ref}}(t)|$ and $|\mathbf{v}(t) - \mathbf{v}_{\mathrm{ref}}(t)|$, are plotted, together with the magnitude of the acceleration compensation command $|\mathbf{a}_{\mathrm{cmp}}(t)|$ and total $\Delta V(t)$ (i.e., reference plus compensation). For clarity, failure cases (i.e., spacecraft crashing on or diverging from Phobos) are not included in the plots, but will be accounted for later on.

From Fig. 16, it can be confirmed that, while all the compensators behave similarly in nominal conditions, the non-robust controller K_0 (Fig. 16a) performs significantly worse than the others under dispersed conditions, particularly in terms of position and velocity error. Additionally, it is clear that, although determined using a reduced set of uncertainties, the WCs identified via μ analysis are generally more challenging than the 2000 MC runs. In fact, these WCs typically have an extremely low probability of occurrence, but could be analytically foreseen, which demonstrates the added-value of μ analysis as a complement to MC verification campaigns.
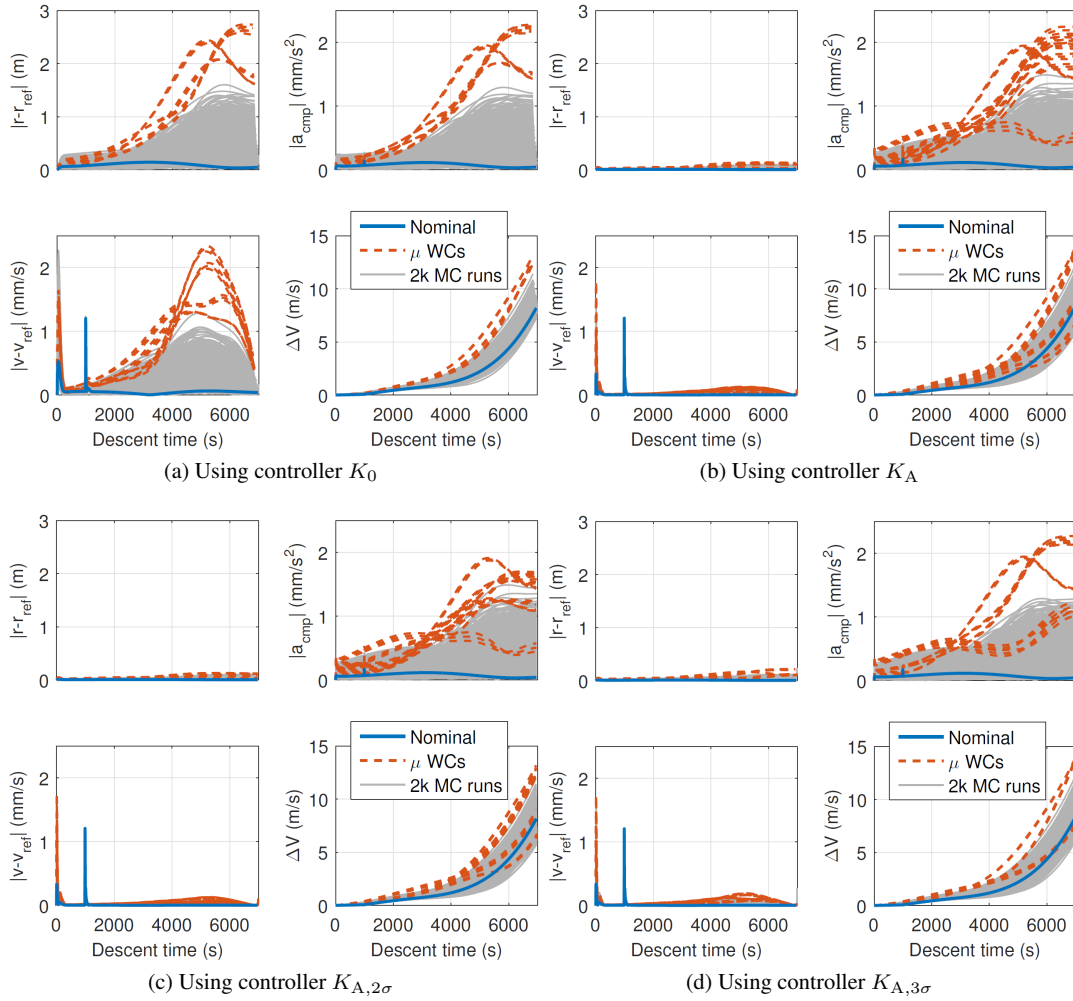
Figure 16. Nonlinear simulation of 2000 MC runs and WCs from μ analysis of RT1

For a detailed comparison, the average and standard deviation of the maximum values encountered during the four MC campaigns are listed in Table II, with the failure ratios of MC and WC simulations.

Table II. Robustness indicators for RT1 using different controllers

	K_0		K_A		$K_{A,2\sigma}$		$K_{A,3\sigma}$	
Max. MC values	Avg.	Std.	Avg.	Std.	Avg.	Std.	Avg.	Std.
$|\mathbf{r} - \mathbf{r}_{\mathrm{ref}}|$ (m)	0.503	0.259	0.029	0.013	0.030	0.014	0.038	0.020
$|\mathbf{v} - \mathbf{v}_{\mathrm{ref}}|$ (mm/s)	1.234	0.130	1.206	0.008	1.208	0.007	1.212	0.013
$|\mathbf{a}_{\mathrm{cmp}}|$ (mm/s^2)	0.411	0.207	0.431	0.208	0.430	0.208	0.429	0.206
ΔV (m/s)	8.793	0.630	8.326	0.966	8.326	0.965	8.329	0.955
MC failures (%)	44.15		0.00		0.00		0.10	
WC failures (%)	33.33		12.50		0.00		16.67	

As anticipated from Fig. 14a, K_0 is not suitable for this scenario because of its higher errors, which are further translated into MC and WC failure ratios of 44.15 and 33.33%. The three robust controllers then showed satisfactory and equivalent distribution indicators, but distinct failure ratios.

With K_A, no failures were obtained from the MC simulations, but 12.5% of the μ WCs resulted into failure. The situation improves with $K_{\mathrm{A},2\sigma}$, which yields no MC or WC failures.

Finally, results using $K_{\mathrm{A},3\sigma}$ became inferior, with failure ratios of 0.1 and 16.37%. Although this last result may seem counter-intuitive, it is actually a fairly common pitfall of robust control algorithms, in which controller performance may become restricted (in the nonlinear case) for an over-conservative robustness specification during the (linear) design phase. It is also important to note that the RP insights from μ analysis only have guaranteed validity for the LTI models adopted, hence the execution of MC tests is fundamental for controller validation in the real, nonlinear case.

Based on Table II, $K_{\mathrm{A},2\sigma}$ is selected as the most suitable controller. To assess its applicability to trajectories other than the one it has been designed for, the same MC uncertainty combinations are executed using $K_{\mathrm{A},2\sigma}$ and trajectories RT2 and RT3 (recall Fig. 3b). The same indicators are now gathered in Table III for these campaigns. Since no detailed μ analysis was performed along RT2 or RT3, the consideration of WCs is not pursued here.

Table III. Robustness indicators for different trajectories using $K_{\mathrm{A},2\sigma}$

	RT1		RT2		RT3	
Max. MC values	Avg.	Std.	Avg.	Std.	Avg.	Std.
$|\mathbf{r} - \mathbf{r}_\mathrm{ref}|$ (m)	0.030	0.014	0.026	0.010	0.026	0.011
$|\mathbf{v} - \mathbf{v}_\mathrm{ref}|$ (mm/s)	1.208	0.007	0.694	0.253	0.661	0.325
$|\mathbf{a}_\mathrm{cmp}|$ (mm/s^2)	0.430	0.208	0.332	0.160	0.326	0.164
ΔV (m/s)	8.326	0.965	8.714	0.831	7.989	0.830
MC failures (%)	0.00		0.00		0.00	

The results of Table III show that $K_{\mathrm{A},2\sigma}$ successfully handles the RT2 and RT3 trajectories with similar performance to RT1 and also without any failure. Since the controller was designed around perturbation variables, the effect of trajectory changes is much reduced in comparison with the considered GH uncertainties. Specific differences among the indicators are mostly related to how demanding is each reference trajectory. These results indicate that accounting for the uncertainties encountered along RT1 alone turned out to be enough for the design of a compensator that is equally able to cope with other trajectories. This conclusion is tightly linked to the previous observation that having a holistic uncertainty specification is not necessarily sufficient (nor required) to achieve a successful robust control design.

6. CONCLUSIONS

This article illustrated the manner gravity uncertainty modelling using linear fractional transformations (LFTs) and structured $\mathcal{H}_\infty$ control synthesis can be effectively employed to design simple yet effective compensators for Space D&L. This in turn paves the way for more reliable planetary landers, as well as guidance approaches with less stringent propellant consumption requirements.

Emphasis was placed on developing a robust control design that accounts for the aforementioned uncertainties and on the suitability of the structured $\mathcal{H}_\infty$ paradigm for the D&L problem. This is mainly because of its ability to synthesise controllers that are valid for multiple plants, as well as the possibility to specify the controller structure and take advantage of industrial legacy knowledge.

Design challenges have also been highlighted, for example on the sensitivity of structured $\mathcal{H}_\infty$ to the choice of initial conditions and control parameters tune.

It was also shown that, more than a robustness assessment tool, μ analysis can be directly employed to support the complete control compensator design process. On the one hand, critical uncertainty combinations can be used to complement and enhance conventional Monte-Carlo verification campaigns and, on the other, closed-loop sensitivities that quantify the impact of each uncertain parameter can be used to support system modelling choices. The practical implications provided by these insights have the potential to span from compensator design to Space D&L mission design.

APPENDIX A. NOMENCLATURE

BCBF	Body-Centred Body-Fixed
SC	Spacecraft
D&L	Descent & Landing
GH	Gravity Harmonics
LFT	Linear Fractional Transformation
LPO	Libration Point Orbit
LTI	Linear Time-Invariant
MC	Monte-Carlo
NP	Nominal Performance
RP	Robust Performance
RS	Robust Stability
RT	Reference Trajectory
SDK	Spacecraft Dynamics & Kinematics
TSA	Testing, Simulation & Assessment
WC	Worst-Case

REFERENCES

1. Falcoz A, Pittet C, Bennani S, Guignard A, Bayart C, Frapard B. Systematic design methods of robust and structured controllers for satellites. *Space Journal* 2015; **7**(3):319–334.
2. Barraclough S, Ratcliffe A, Buchwald R, Chapuy M, Garland M, Rebuffat D. Phootprint: A European Phobos Sample Return Mission. *The 11th International Planetary Probe Workshop*, Passadena, CA, 2014.
3. Joffre E, Zamaro M, Silva N, Marcos A, Simplício P, Richardson B. Landing on Small Bodies Trajectory Design, Robust Nonlinear Guidance and Control. *The 27th AAS/AIAA Spacefilght Mechanics Meeting*, San Antonio, TX, 2017.
4. Simplício P, Marcos A, Joffre E, Zamaro M, Silva N. Parameterised Laws for Robust Guidance and Control of Planetary Landers. *The 4th CEAS EuroGNC Conference*, Warsaw, Poland, 2017.
5. Simplício P, Marcos A, Joffre E, Zamaro M, Silva N. A Systematic Performance-oriented Tuning for Space Exploration Descent & Landing Guidance. *The 7th European Conference for Aeronautics and Space Sciences*, Milan, Italy, 2017.
6. Schmit N. Robust Control of Satellite with Fuel Slosh Instabilities. *The 60th International Astronautical Congress*, Daejeon, South Korea, 2009.
7. Charbonnel C. $\mathcal{H}_\infty$ Controller Design and μ-Analysis: Powerful Tools for Flexible Satellite Attitude Control. *The 2010 AIAA Guidance, Navigation, and Control Conference*, Toronto, Canada, 2010.

8. Watt M, Yu M, Falcoz A, Singh P, Warren C. BIOMASS Normal Mode AOCS: Classical versus Robust Design. *The 8th International ESA Conference on Guidance, Navigation and Control Systems*, Karlovy Vary, Czech Republic, 2011.

9. Gahinet P, Apkarian P. Structured $\mathcal{H}_\infty$ Synthesis in MATLAB. *The 18th IFAC World Congress*, Milan, Italy, 2011.

10. Apkarian P, Dao M, Noll D. Parametric robust structured control design. *Transactions on Automatic Control* 2015; **60**(7):1857–1869.

11. Zamaro M, Biggs J. Natural motion around the Martian moon Phobos: the dynamical substitutes of the Libration Point Orbits in an elliptic three-body problem with gravity harmonics. *Celestial Mechanics and Dynamical Astronomy* 2015; **122**(3):263–302.

12. Hawkins M, Guo Y, Wie B. Spacecraft Guidance Algorithms for Asteroid Intercept and Rendezvous Missions. *International Journal of Aeronautical and Space Sciences* 2012; **13**(2):154–169.

13. Battin R. *An Introduction to the Mathematics and Methods of Astrodynamics*. 1st edn., AIAA Education Series, 1987.

14. Doyle J, Packard A, Zhou K. Review of LFTs, LMIs and μ. *The 30th IEEE Conference on Decision and Control*, Brighton, UK, 1991.

15. Zhou K, Doyle J, Glover K. *Robust and Optimal Control*. 1st edn., Prentice-Hall, 1995.

16. Balas G, Chiang R, Packard A, Safonov M. *MATLAB Robust Control Toolbox User's Guide*. 3rd edn., MathWorks, 2006.

17. Roos C, Hardier G, Biannic JM. Polynomial and rational approximation with the APRICOT Library of the SMAC toolbox. *The 2014 IEEE Conference on Control Applications*, Antibes, France, 2014.

18. Pignié G. Ariane 5 and Ariane 5 evolution GN&C overview. *The 2nd World Space Congress*, Houston, TX, 2002.

19. Bourdon J, Delpy P, Ganet M, Quinquis I, Ankersen F. Application of $\mathcal{H}_\infty$ Design on ATV Control Loop During the Rendezvous Phase. *The 5th ESA International Conference on Spacecraft Guidance, Navigation and Control Systems*, Frascati, Italy, 2002.

20. Apkarian P, Noll D. Nonsmooth $\mathcal{H}_\infty$ Synthesis. *Transactions on Automatic Control* 2006; **51**(1):71–86.

21. Saussié D, Barbès Q, Bérard C. Self-Scheduled and Structured $\mathcal{H}_\infty$ Synthesis: a Launch Vehicle Application. *The 2013 American Control Conference*, Washington, DC, 2013.

22. Lhachemi H, Saussié D, Zhu G. A Robust and Self-Scheduled Longitudinal Flight Control System: a Muli-Model and Structured $\mathcal{H}_\infty$ Approach. *The AIAA SciTech 2014 Forum*, National Harbor, MD, 2014.

23. Pittet C, Prieur P. Structured accelero-stellar estimator for MICROSCOPE drag-free mission. *The 3rd CEAS EuroGNC Conference*, Toulouse, France, 2015.

24. Marcos A, Sato M. Flight Testing of a Structured H-infinity Controller: a EU-Japan Collaborative Experience. *The 1st IEEE Conference on Control Technology and Applications*, Kohala Coast, HI (pending acceptance), 2017.

25. Stein G, Doyle J. Beyond Singular Values and Loop Shapes. *Journal of Guidance, Control, and Dynamics* 1991; **14**(1):5–16.

26. Marcos A, Bates D, Postlethwaite I. Control oriented uncertainty modelling using μ sensitivities and skewed μ analysis tools. *The 44th IEEE Conference on Decision and Control*, Seville, Spain, 2005.

27. Simplício P, Bennani S, Marcos A, Roux C, Lefort X. Structured Singular-Value Analysis of the Vega Launcher in Atmospheric Flight. *Journal of Guidance, Control, and Dynamics* 2016; **39**(6):1342–1355.

28. Biannic JM, Burlion L, Demourant F, Ferreres G, Hardier G, Loquen T, Roos C. The SMAC Toolbox: a collection of libraries for Systems Modeling, Analysis and Control. Online available at http://w3.onera.fr/smac/, 2016.