

A Systematic Robust Control Design and Verification Approach for Drag-free Space Missions

Original

A Systematic Robust Control Design and Verification Approach for Drag-free Space Missions / Marchioli, V., Chiriaco', M., Basile, F., Boggio, M., Pagone, M., Novara, C.. - In: AEROSPACE SCIENCE AND TECHNOLOGY. - ISSN 1270-9638. - (2026). [10.1016/j.ast.2026.114014]

Availability:

This version is available at: 11583/3016452 since: 2026-10-09T09:19:21Z

Publisher:

Elsevier

Published

DOI:10.1016/j.ast.2026.114014

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright


(Article begins on next page)

Journal Pre-proof

A Systematic Robust Control Design and Verification Approach for Drag-free Space Missions

Vinicius Marchioli, Marco Chiriaco, Francesco Basile, Mattia Boggio, Michele Pagone, Carlo Novara

PII: S1270-9638(26)02390-4
DOI: <https://doi.org/10.1016/j.ast.2026.114014>
Reference: AESCTE 114014



To appear in: *Aerospace Science and Technology*

Received date: 23 July 2026
Revised date: 20 September 2026
Accepted date: 5 October 2026

Please cite this article as: Vinicius Marchioli, Marco Chiriaco, Francesco Basile, Mattia Boggio, Michele Pagone, Carlo Novara, A Systematic Robust Control Design and Verification Approach for Drag-free Space Missions, *Aerospace Science and Technology* (2026), doi: <https://doi.org/10.1016/j.ast.2026.114014>

This is a PDF of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability. This version will undergo additional copyediting, typesetting and review before it is published in its final form. As such, this version is no longer the Accepted Manuscript, but it is not yet the definitive Version of Record; we are providing this early version to give early visibility of the article. Please note that Elsevier's sharing policy for the Published Journal Article applies to this version, see: <https://www.elsevier.com/about/policies-and-standards/sharing#4-published-journal-article>. Please also note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2026 Published by Elsevier Masson SAS.

Highlights

- A systematic approach for robust control design and verification approach to drag-free space missions is presented.
- The design complexity is progressively increased from decoupled SISO synthesis to MIMO robust refinement.
- Well-established H_∞ and μ -analysis tools are nested within a unified hierarchical design and verification workflow.
- The applicability and generality of the approach are demonstrated using LISA and NGGM space missions.

A Systematic Robust Control Design and Verification Approach for Drag-free Space Missions

Vinicius Marchioli^a, Marco Chiriaco^a, Francesco Basile^b, Mattia Boggio^a,
Michele Pagone^a, Carlo Novara^a

^a*Department of Electronics and Telecommunications, Politecnico di Torino, Corso Duca degli Abruzzi, 24, Torino, 10129, Italy*

^b*Thales Alenia Space Italy, Strada Antica di Collegno, 253, Turin, 10146, Italy*

Abstract

A systematic and iterative framework for robust control design and verification, tailored to critical drag-free space missions, is presented. The main contribution of the work consists of integrating well-established $\mathcal{H}_\infty$ and μ -based tools within a unified hierarchical design and verification workflow specifically devised for drag-free applications, in which the design complexity is progressively increased from decoupled single-input single-output synthesis to multi-input multi-output robust refinement only when required by the verification results, thereby limiting unnecessary controller complexity, while systematically addressing mission-level stability and performance requirements. By integrating linear robustness analysis with parametric uncertainty considerations, the framework provides a systematic procedure to assess drag-free compliance with mission-critical constraints and is applied to the robust control design and analysis of the European Space Agency missions LISA and NGGM. Numerical simulations confirm the effectiveness of the approach, highlighting its adaptability and potential application to future advanced drag-free space missions.

¹*Email addresses:* `vinicius.marchioli@polito.it` (Vinicius Marchioli),
`marco.chiriaco@polito.it` (Marco Chiriaco),
`francesco.basile@thalesaleniaspace.com` (Francesco Basile),
`mattia.boggio@polito.it` (Mattia Boggio),
`michele.pagone@polito.it` (Michele Pagone),
`carlo.novara@polito.it` (Carlo Novara).

Keywords: Laser interferometry, drag-free attitude control, robust stability, μ -analysis, gravimetric missions.

1. Introduction

In the last four decades, the European Space Agency (ESA) has devoted a significant part of its scientific program to ambitious space missions within *Horizon 2000* (1984–2018) and *Cosmic Vision* (2015–present), aimed at investigating fundamental questions in astrophysics, cosmology, and Earth science. Relevant examples include the Gravity Field and Steady-State Ocean Circulation Explorer (GOCE) mission [1], launched in March 2009, and the Next Generation Gravity Mission (NGGM) [2] [3], planned for 2032 and aimed at long-term monitoring of temporal variations of the Earth’s gravity field. Another prominent example is the Laser Interferometer Space Antenna (LISA) mission, expected to be launched in 2034, aimed at detecting gravitational waves through precision laser interferometry [4] [5] [6]. A common feature of these missions is the complexity of spacecraft (SC) dynamics, which are affected by multiple disturbance sources that are difficult to predict and model accurately, making the design of control systems capable of satisfying stringent orbital and attitude requirements a key challenge [7] [8].

For example, in LISA, gravitational-wave detection relies on measuring the relative distance between test masses (TMs) belonging to different spacecraft of the constellation, and achieving the TM free-fall condition [9] requires minimizing the effects of non-gravitational forces, such as solar radiation pressure and residual spacecraft accelerations [10] [11].

In this paper, we focus on Drag-Free Attitude Control System (DFACS) design and analysis. DFACS aims to minimize non-gravitational disturbances in order to preserve the quality of the scientific measurements; since complete rejection of external, internal, sensing, and actuation disturbances is not physically achievable, the drag-free controller must attenuate their effects within a prescribed frequency bandwidth. DFACS is therefore particularly relevant to high-precision astrophysics, fundamental physics, and Earth observation missions, where small disturbances can significantly affect scientific measurements, and missions such as LISA and GRACE (Gravity Recovery and Climate Experiment) [12] rely on drag-free control to preserve the quality of the scientific observables.

In recent years, several drag-free control design concepts have been proposed. Embedded Model Control [13] [14] estimates disturbances affecting

the plant and compensates for disturbances and unmodeled dynamics, while a two-stage design technique guaranteeing robust performance is proposed in [15] and [16], where bounds on the closed-loop responses are first derived and decoupled single-input single-output (SISO) controllers are subsequently obtained through optimization. Direct space-mission applications include the $\mathcal{H}_\infty/\mu$ -synthesis approach in [17], which develops a multi-input multi-output (MIMO) robust structured controller, and the $\mathcal{H}_\infty$ -based DFACS proposed for LISA in [10], based on constrained SISO decoupling with stability guarantees. Other approaches include genetic-algorithm-based LQR tuning [18] and partial state-feedback adaptive control [19]. Nevertheless, many of these studies focus on mission-specific design solutions, while a systematic DFACS design and verification framework is still missing.

In this paper, we propose a systematic approach for DFACS design and verification for drag-free space missions. The framework combines structured $\mathcal{H}_\infty$ synthesis, μ -analysis, and optional μ -synthesis within a progressive-complexity design and verification workflow. The methodology addresses both SISO and MIMO control synthesis, followed by stability, performance, and robustness verification. A low-complexity SISO architecture is retained whenever it remains compliant after reconstruction of the complete uncertain MIMO model, while MIMO synthesis is introduced only when required by the robustness assessment. The control design aims to satisfy scientific requirements in the frequency domain together with system-level constraints, with the $\mathcal{H}_\infty$ framework adopted to address performance and stability in the presence of disturbances and uncertainties, while robust stability and performance are verified through μ -analysis (see, e.g., [20] [21]), which also supports subsequent controller refinement through μ -synthesis.

Hence, the main contribution of this work does not lie in a novel robust control and analysis mechanism, but rather in the integration of established $\mathcal{H}_\infty$ and μ -based tools within a unified hierarchical and iterative framework specifically devised for drag-free applications. The proposed strategy starts from low-complexity decoupled SISO controllers and increases the design complexity only when the complete MIMO robustness verification identifies critical interactions or insufficient robustness margins, with μ -analysis acting both as a verification and diagnostic tool to guide local SISO retuning or MIMO robust refinement when required.

A further contribution lies in the mission-independent workflow structure. Mission-specific information is introduced through the plant model, performance requirements, disturbance spectral descriptions, weighting functions,

and structured uncertainty blocks, while the overall design and verification sequence remains unchanged. Unlike conventional verification approaches that often rely on Monte Carlo campaigns, the proposed methodology primarily exploits the linear design environment for robustness verification, while nonlinear or statistically driven simulations are retained as complementary analyses when required. The applicability and generality of the approach are demonstrated using LISA and NGGM as representative benchmarks, characterized by different scientific objectives, controlled variables, sensing and actuation architectures, operating frequency bands, and uncertainty models.

The paper is organized as follows. Section 2 introduces the DFACS problem, while Section 3 describes the proposed design and verification framework; Sections 4–5 present the LISA and NGGM case studies, respectively, and conclusions are drawn in Section 6.

2. Problem Preliminaries

Consider a continuous-time nonlinear system affected – in an affine fashion – by an exogenous, unknown, and bounded disturbance

$$\dot{x}(t) = f(x(t), u(t)) + B_w w(t), \quad t \geq 0, \quad (1)$$

where $x \in \mathbb{R}^{n_x}$, $u \in \mathbb{R}^{n_u}$, and $w \in \mathbb{R}^{n_w}$ denote the state, input, and disturbance at time t , respectively.

Consider a disturbance-free equilibrium pair $(\bar{x}, \bar{u})$ satisfying $f(\bar{x}, \bar{u}) = 0$, around which the linearized version of (1), including measured and performance outputs, is given by

$$\begin{aligned} \dot{x}(t) &= Ax(t) + B_u u(t) + B_w w(t), \\ v(t) &= C_v x(t) + D_{vu} u(t) + D_{vw} w(t), \\ z(t) &= C_z x(t) + D_{zu} u(t) + D_{zw} w(t), \end{aligned} \quad (2)$$

where $v \in \mathbb{R}^{n_v}$ denotes the vector of measured outputs and $z \in \mathbb{R}^{n_z}$ collects the performance variables. The state vector x represents the internal plant dynamics and is not directly available to the controller.

Assumption 1. *We assume that the pair (A, B_u) is stabilizable, the pair (A, C_v) is detectable, and the plant-controller interconnection is well posed.*

The main goal of the DFACS is to attenuate the effects of non-gravitational forces on a SC, ensuring the so-called *free-fall condition*, which is critical for

high-precision space missions requiring extreme accuracy in position and attitude control. For example, in LISA, gravitational-wave detection relies on measuring the relative distance between TMs belonging to different spacecraft of the constellation; in this context, the DFACS must attenuate disturbances within the measurement bandwidth so that the residual attitude and position jitters satisfy stringent requirements [6] [10] [22]. Since complete rejection of external, internal, sensing, and actuation disturbances is not physically achievable, the drag-free controller must minimize their effects within a prescribed bandwidth.

Definition 1. *Consider a closed-loop system composed of a plant and a controller connected in feedback and subject to external disturbances characterized by given spectral densities. Let Y be a vector of output variables of interest, $S_Y^{1/2}(\nu)$ the vector containing their amplitude spectral densities (ASDs), $[\underline{\nu}, \bar{\nu}]$ a frequency bandwidth of interest, and $\bar{S}^{1/2}(\nu)$ a vector of desired ASD requirement bounds. The closed-loop system is drag-free in the prescribed bandwidth if*

$$S_Y^{1/2}(\nu) \leq \bar{S}^{1/2}(\nu), \quad \forall \nu \in [\underline{\nu}, \bar{\nu}], \quad (3)$$

where the inequality is element-wise. A controller satisfying this condition is called a drag-free controller.

We introduce the linear fractional transformation (LFT) of the system under investigation [23]. Let $P(s)$ denote the generalized LTI plant associated with (2), and let $K(s)$ be a stabilizing dynamic output-feedback controller. Given the exogenous input w , the control input u , the measured output v , and the performance output z (see Fig. 1), the input-output relations are

$$\begin{bmatrix} z \\ v \end{bmatrix} = \begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix} \begin{bmatrix} w \\ u \end{bmatrix}, \quad u = K(s)v. \quad (4)$$

The controller $K(s)$ is a dynamic output-feedback controller with the state-space realization

$$K : \begin{cases} \dot{x}_K(t) = A_K x_K(t) + B_K v(t), \\ u(t) = C_K x_K(t) + D_K v(t), \end{cases} \quad (5)$$

where x_K denotes the controller internal state. Therefore, the controller operates on the measured output vector v , rather than on the full plant




Figure 1: General lower-LFT interconnection between the generalized plant $P(s)$ and the dynamic output-feedback controller $K(s)$.

state, with v constructed in the mission-specific applications from the sensing channels actually available to the control system.

From (4), the lower linear fractional transformation is

$$F(P, K) = P_{11} + P_{12}K(s)(I - P_{22}K(s))^{-1}P_{21}, \quad (6)$$

so that $z = F(P, K)w$.

3. Control design and verification methodology

A general overview of the complete DFACS design and verification approach is illustrated in Fig. 2. The proposed methodology is divided into multiple stages addressing both SISO and MIMO control synthesis, followed by stability, performance, and robustness verification.

A distinctive feature of the proposed methodology lies in the decision logic connecting the different design and verification stages. Rather than directly adopting a full MIMO robust synthesis, the framework follows a progressive-complexity strategy, in which low-order decoupled SISO controllers are first designed and subsequently assessed within the complete uncertain MIMO environment. The SISO architecture is retained whenever the required stability and performance levels are preserved, whereas MIMO refinement is introduced only when the robustness analysis identifies critical multivariable interactions or insufficient robustness margins.

Since the first design level deals with linear controllers, system (1) is linearized around the equilibrium pair $(\bar{x}, \bar{u})$, after which a decoupled approx-




Figure 2: Closed-loop general scheme for the systematic design.

imation is introduced, allowing the adoption of a preliminary SISO control architecture (see, e.g., the LISA case in [10]).

Remark 1. *At this stage, the decoupling is limited to the controller synthesis to obtain a low-order baseline controller by neglecting cross-axis interactions, assumed to be weaker than the dominant direct input-output channels. These interactions are reintroduced at the subsequent MIMO robustness verification stage as structured uncertainty blocks; if the resulting μ indicators remain below unity, the SISO architecture is retained, whereas the critical channels are explicitly included in the MIMO synthesis and the controller is refined accordingly when the robustness conditions are not satisfied.*

The control design aims to satisfy requirements associated with multiple degrees of freedom (DoFs), including science objectives within a prescribed frequency bandwidth together with system-level constraints. This motivates the use of an $\mathcal{H}_\infty$ approach, due to its ability to address performance and stability requirements in the presence of uncertainties and disturbances by minimizing the worst-case gain from disturbance inputs to controlled outputs. The $\mathcal{H}_\infty$ design procedure for DFACS is described in [10] [21] and omitted here for brevity, while the overall control design and verification procedure consists of the following steps.

1. *SISO level synthesis:* The nonlinear system is linearized and decomposed into decoupled SISO loops, after which low-order structured $\mathcal{H}_\infty$ -based controllers are designed. Hard constraints are used for design

requirements, while soft constraints shape secondary performance objectives, such as the transient response.

2. *Linear robustness verification:* A robustness evaluation is conducted by applying μ -analysis to the overall MIMO plant with the combined controller structure, and structured singular-value bounds are computed to assess robust stability and robust performance with respect to the modeled uncertainties. The resulting analysis determines whether further controller refinements are necessary.
3. *MIMO level synthesis:* If the SISO controllers fail to meet performance or stability requirements in the multichannel configuration, MIMO optimization techniques, including μ -synthesis, are employed [24], thereby refining problematic interactions and improving robustness by explicitly addressing critical multivariable couplings.
4. *Nonlinear robustness verification:* Requirement satisfaction is verified through long-duration simulations, assessing whether neglected nonlinearities and representative worst-case disturbances degrade DFACS performance.
5. *Closed-loop iteration and termination:* The previous steps are repeated as needed, with adjustments at the SISO or MIMO level. Supporting analyses, such as noise breakdown studies, are used to identify the dominant contributors to performance degradation.

Remark 2. *The proposed iterative workflow is meant as an engineering design-and-verification process rather than a certificate of formal global convergence. A controller is accepted only when the nominal requirements, robust stability and robust performance conditions, and the applicable verification criteria are satisfied; otherwise, further SISO or MIMO refinements are performed. The process is terminated without acceptance when the remaining issues are associated with physical sensing or actuation limitations that cannot be mitigated through controller redesign, or when further iterations do not provide a significant benefit compared with the increase in complexity.*

3.1. SISO level synthesis

The $\mathcal{H}_\infty$ control problem aims to achieve a closed-loop transfer function with the minimum $\mathcal{H}_\infty$ norm; however, the resulting stabilizing controller $K(s)$ often has a very high order (similar to that of $P(s)$). In aerospace applications, low-order controllers are preferred to reduce computational costs, simplify the synthesis process and accelerate the design procedure [25].

Since space control problems, including DFACS applications, often exhibit linearized dynamics dominated by double-integrator behavior, we adopt the following structured controller $K(s)$:

$$K(s) = K_p \cdot K_{pl}(s) \cdot K_{lp}(s) \cdot K_{DC}(s). \quad (7)$$

Consistently with the dynamic output-feedback formulation introduced in Section 2, the controller in (7) is a dynamic compensator.

In (7), K_p is a proportional gain, $K_{pl}(s) = (s + z_1)/(s + p_1)$ is a phase-lead contribution, $K_{lp}(s) = \omega_n/(s + \omega_n)$ is a fixed first-order low-pass filter, and $K_{DC}(s) = [(s + \omega_{n1})/(s + \omega_{n2})]^2$ is a second-order DC-shaping filter. Here, ω_n is the low-pass cutoff angular frequency, while ω_{n1} and ω_{n2} are the zero and pole angular frequencies of $K_{DC}(s)$, respectively; for $\omega_{n1} < \omega_{n2}$, $K_{DC}(s)$ exhibits a high-pass shelving behavior, attenuating low-frequency and DC components while approaching unit gain at higher frequencies.

Each SISO controller follows the depicted structure. The algorithm addresses the same $\mathcal{H}_\infty$ performance objective as the classical unstructured synthesis, while restricting the controller to a prescribed tunable structure. Mandatory requirements are enforced as hard constraints, also through weighted output functions, while secondary control objectives are introduced as soft constraints [10].

The effectiveness of the SISO design level is evaluated according to the following criteria: i) feasibility of the design; ii) the resulting cut-off frequency must be at least twice the unstable cut-off frequency associated with negative stiffness phenomena; iii) fulfillment of the single-channel requirements; iv) a phase margin of ≥ 30 degrees and a gain margin of ≥ 6 dB are expected according to [26]; and v) acceptable overall behavior in terms of nonlinearities, inter-channel couplings, and sensing and actuation cross-talks. A design is retained when the SISO loops successfully pass the subsequent robustness verification.

3.2. Linear robustness verification

Linear robustness verification is performed through μ -analysis to assess the worst-case closed-loop behavior with respect to the modeled structured uncertainties, using a generalized plant that includes the plant, sensing, actuation, and measurement dynamics together with the internal dynamics of the implemented controllers. The corresponding LFT model, illustrated in Fig. 3, incorporates the uncertainty block Δ .




Figure 3: LFT interconnection of the system for μ -analysis. Δ represents the uncertainty block, P is the plant and K is the controller.

The closed-loop system defined by the feedback connection between $P(s)$ and $K(s)$ can be written as

$$M = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix},$$

with inputs $\eta \in \mathbb{R}^{n_\eta}$ and $w \in \mathbb{R}^{n_w}$, and outputs $y \in \mathbb{R}^{n_y}$ and $z \in \mathbb{R}^{n_z}$, where (y, η) define the uncertainty interconnection, and (w, z) define the exogenous-input/performance channels.

Since $K(s)$ is dynamic, the nominal closed-loop state is the augmented vector $x_{CL} = [x_P^\top x_K^\top]^\top$, including both plant and controller states:

$$M : \begin{cases} \dot{x}_{CL} &= A_{CL}x_{CL} + B_\eta\eta + B_w w, \\ y &= C_y x_{CL} + D_{y\eta}\eta + D_{yw} w, \\ z &= C_z x_{CL} + D_{z\eta}\eta + D_{zw} w. \end{cases} \quad (8)$$

Here, A_{CL} is the augmented state matrix obtained from the dynamic output-feedback interconnection between $P(s)$ and $K(s)$, while the uncertainty block Δ collects the normalized structured uncertainties considered in the analysis, including real parametric uncertainties, repeated blocks, and cross-coupling terms (as specified for each mission case study).

Definition 2. Given two complex matrices M and Δ , we define the structured singular value as

$$\mu_\Delta(M) \doteq \frac{1}{\bar{\sigma}(\Delta_{\min})},$$

where

$$\Delta_{\min} \doteq \arg \min_{\Delta} \{ \bar{\sigma}(\Delta) : \Delta \in S_{\Delta}, \det(I - M\Delta) = 0 \},$$

and $\bar{\sigma}(\cdot)$ denotes the largest singular value; if no such Δ exists, then $\mu_{\Delta}(M) := 0$.

According to [27] and [21], the robust stability criterion holds if and only if

$$\mu_{\Delta}^s(M) \doteq \sup_{\omega \in \mathbb{R}^+} \mu_{\Delta}(M_{y\eta}(j\omega)) < 1, \quad (9)$$

provided that the nominal interconnection is internally stable and Δ belongs to the prescribed normalized structured uncertainty set.

Upper bound on $\mu_{\Delta}(M)$ satisfy

$$\max_{U \in U} \rho(UM) \leq \mu_{\Delta}(M) \leq \inf_{D \in D} \bar{\sigma}(DMD^{-1}), \quad (10)$$

where U and D denote the admissible matrices associated with the prescribed uncertainty structure [27], $\bar{\sigma}(\cdot)$ is the largest singular value, and $\rho(\cdot)$ is the spectral radius.

Beyond robust stability, robust performance is assessed by augmenting the physical uncertainty structure Δ with a fictitious performance block associated with the disturbance-to-performance channel $w \rightarrow z$. The resulting analysis verifies whether the prescribed performance requirements remain satisfied for all admissible uncertainty realizations.

According to Sec. 6.8.2 of the ESA ECSS Control Performance Handbook [26], the uncertainty set should be regarded as representative and non-exhaustive, based on the dominant physical properties and uncertainty sources of the system.

In this work, the uncertainty domain includes the parameters expected to have the largest impact on the closed-loop behavior: spacecraft and TM inertial properties, TM stiffness, sensing and actuation cross-talks, relevant system dynamics, and system delays.

In summary, robust stability is assessed on the uncertainty feedback interconnection, whereas robust performance also includes the relevant disturbance and performance channels, with the corresponding peak μ upper bound required to remain below unity over the frequency range of interest in both cases.

Additionally, a noise breakdown analysis can identify the main performance contributors and support subsequent controller refinement; if the robustness criteria are not met, individual SISO controllers can be retuned when feasible, whereas MIMO μ -synthesis may be considered otherwise.

3.3. MIMO level synthesis

Specific channels that struggle to fulfill stability or performance robustness requirements may be tuned using a MIMO μ -synthesis approach [28], which increases the complexity of the optimization problem but allows multiple critical channels to be addressed simultaneously. The iterative process starts with a model augmentation, incorporating additional constraints and structured parametric uncertainties, and extending the design to selected multivariable subsystems, after which μ -synthesis is used to perform the D–K iteration. The D–K iteration is stopped when it does not further reduce the peak structured singular value, or when the corresponding increase in controller complexity does not offer any robustness improvement.

Alternatively, a structured MIMO design may be adopted to balance the simplicity of the SISO design and the conservativeness of μ -synthesis. Instead of tackling the entire system at once, a subset of input and output channels is selected based on sensitivity analysis, couplings, or their contribution to critical performance metrics. This reduces the problem dimensionality while preserving the possibility of applying linear design techniques.

At this level, compliance follows the same logic as for the SISO approach, with each iteration evaluating stability margins, performance robustness, and interaction effects among the selected channels. Compliance is considered achieved when all critical channels satisfy their performance and robustness criteria without compromising overall system stability and when the nonlinear robustness verification is successful.

3.4. Nonlinear robustness verification

The nonlinear robustness verification is carried out at MIMO level through extended simulations using the nonlinear high-fidelity model of the plant and considering representative worst-case disturbance scenarios.

These simulations verify that the spectral densities of the residual drag-free acceleration/jitter remain below the prescribed frequency-domain requirements within the mission bandwidth, while relevant actuator magnitude saturation and rate limitations are assessed directly in the time domain,

since frequency-domain constraints alone do not guarantee that such limits are never exceeded under large disturbances.

The simulations span several mission days to assess whether combinations of worst-case disturbances and nonlinearities neglected during the control design and linear verification degrade the closed-loop robust performance.

4. Case Study I: The LISA Mission

The LISA mission is one of ESA's most complex and ambitious long-term missions, designed to detect and study low-frequency gravitational waves in the range of 0.1 mHz to 1 Hz [29] [30]. It comprises a constellation of three satellites separated by 2.5 million kilometers, arranged in a triangular formation, and connected through bidirectional laser links to perform precision interferometry. Each satellite houses free-falling TMs (see Fig. 4) that must be protected from non-gravitational disturbances, categorizing LISA as a drag-free mission.

Unlike current ground-based observatories like Virgo [31] and LIGO [32], LISA's space-borne platform is planned to operate in the quieter environment of space, extending gravitational-wave observations toward much lower frequencies and enabling the observation of sources such as supermassive black hole mergers, galactic binaries, and compact binary systems within the Milky Way. LISA may also detect relic gravitational waves from the early Universe and other currently unexplored sources.

The complete set of equations governing the nonlinear dynamics of LISA is presented in [6], while the linearized model is reported in [21]. For brevity, these equations are not reproduced here.

4.1. LISA science-mode overview

The primary objective of LISA's science mode is to maintain the two TMs at the center of the electrostatic housing cage, while preserving their free-fall condition along the interferometer sensitive axis. This configuration allows non-gravitational disturbances along the interferometry direction to be minimized, where displacements on the picometer scale shall be detected. Simultaneously, the SC and Moving Optical Sub-Assemblies (MOSAs) are commanded to follow the triangular constellation formation (see Fig. 4) through multiple control loops integrating payload and platform sensors and actuators within the DFACS. The controller must therefore satisfy stringent performance requirements while compensating for parametric uncertainties,




Figure 4: LISA configuration overview.

sensor/actuator noise, and external forces and torques affecting the spacecraft and the overall measurement dynamics.

Relevant perturbations include solar radiation pressure, self-gravity, TM stiffness variations, environmental and impulsive disturbances, and sensing and actuation noise.

Table A.7 in the Appendix summarizes the primary performance requirements for LISA's science mode. While drag-free missions typically constrain residual acceleration along the experimental axis, in LISA the top-level requirement is expressed as relative jitter for each subsystem DoF in the frequency domain.

The science-mode DFACS coordinates four main subsystems, resulting in a total of 18 controlled DoFs for each LISA spacecraft.

1. **Drag-Free Loop.** The spacecraft controls its three translational degrees of freedom through the Micro-Propulsion System (MPS) to maintain the TMs in free fall. Measurements are provided by interferometric readouts along the sensitive axis and by the Gravitational Reference Sensor (GRS), which monitors the relative position and motion of the TMs with respect to the spacecraft, enabling the estimation of the relative motion required for drag-free control.
2. **Test Mass Suspension Loop.** The TM degrees of freedom — $TM_{\theta,1}$, $TM_{\theta,2}$, $TM_{\eta,1}$, $TM_{\eta,2}$, $TM_{\phi,1}$, $TM_{\phi,2}$, $TM_{y,1}$, $TM_{y,2}$, $TM_{z,1}$, and $TM_{z,2}$ — define the suspension loop, in which each TM must remain centered within its electrode housing. Here, θ , η , and ϕ denote small rotations about three orthogonal spacecraft-fixed axes, while y and z are

translational displacements orthogonal to the main sensitive axis. A short-arm Differential Wavefront Sensing (DWS) system measures η and ϕ , while the GRS provides the remaining translational and rotational measurements; based on these measurements, the electrostatic suspension system applies the required forces and torques to keep the TMs centered.

3. **Spacecraft Attitude Loop.** The spacecraft attitude is described by the three small-angle rotations SC_θ , SC_η , and SC_ϕ with respect to the constellation reference frame. These DoFs are controlled through cold-gas micro-propulsion torques, while fine attitude measurements are provided by the DWS system.
4. **Pointing Loop.** The two optical assemblies must be precisely oriented to maintain the required optical alignment within the triangular constellation, which is achieved by controlling $MOSA_{\phi,1}$ and $MOSA_{\phi,2}$ around the nominal 60° telescope separation angle. These mechanisms ensure the required long-term pointing stability in coordination with the spacecraft attitude control loop.

4.2. Design results

To evaluate the methodology introduced in Section 3, the TM suspension control loop of a single TM along y and z is designed and assessed for the LISA science mode, due to its stringent requirements and sensitivity to noise.

Remark 3. *Note that, while the x axis is traditionally the primary interferometer axis in drag-free missions, ESA removed the corresponding acceleration requirement from the DFACS requirements baseline, retaining only y and z position requirements as binding constraints. The x axis requirement may nonetheless be reinstated in future mission iterations and addressed within the same methodological framework.*

The SISO level design is based on the closed-loop architecture depicted in the following Fig. 5.




Figure 5: Generic closed-loop SISO diagram with appended weighting functions.

A structured controller, K , is synthesized according to Eq. (7), with the optimization considering the control effort u , disturbance input d , measurement noise n , and regulation error e . Hard constraints are imposed on the channels from n and d to u , while soft constraints are used to shape the complementary sensitivity function. To avoid excessively slow controller dynamics, a minimum decay-rate constraint corresponding to 0.1 Hz is imposed, thereby excluding very slow poles from the resulting controller.

The hard constraints are shown in Figs. 6a and 6b, where the closed-loop responses remain below the prescribed bounds reported in Table A.7.




Figure 6: Hard constraints for the TM suspension loop.

Similarly, the soft constraints are shown in Figs. 7a and 7b, confirming consistency with the prescribed transient-performance objectives.





Figure 7: Soft constraints for TM suspension loop.

Figs. 8 and 9 provide a source-by-source spectral breakdown of the output channels, with each disturbance and sensing-noise input activated independently. The corresponding spectra follow the prescribed masks in Table A.7, which are treated as frequency-domain envelopes without assuming Gaussian disturbance processes. The $\mathcal{H}_\infty$ synthesis constrains the induced closed-loop gains independently of any specific disturbance probability distribution, while noise correlations are not included in this decomposition and would require a joint cross-spectral description.

Figure 8: Contribution to the output channel for the TM suspension loop in y .

The final low-order structured controller, $K(s)$, satisfies the single-channel requirements and provides a sufficiently large cut-off frequency to mitigate




Figure 9: Contribution to the output channel for the TM suspension loop in z .

negative-stiffness effects. It is then assessed through the subsequent MIMO robustness verification.

4.3. Robustness Verification Results

The uncertainty ranges considered in the μ -analysis are summarized in the Appendix (Table A.8) and modeled as structured real parametric perturbations around their nominal values. Note that, such values are provided by ESA according to the current development and state-of-the-work of the LISA mission. Such values are likely to change in the future, depending on advancements and developments of the LISA SCs design. Robust stability and robust performance are assessed according to the framework introduced in Section 3.2.

As a preliminary robustness indicator, classical gain and phase margins are evaluated during the science phase and reported in Table 1, with the obtained phase margins exceeding 30° .

Table 1: Nominal DFACS stability margins.

Loop	Axis	Phase Margin [deg]	Gain Margin [dB]
TM	y	42.8	-16.7
	z	35.1	-12.4

Subsequently, a structured singular value μ -analysis is performed to quan-

tify robustness under the considered uncertainties. The robust-performance channels include the TM y and z translational jitter and GRS actuation noise (u_y, u_z) , with the corresponding GRS actuation and sensor-noise spectra defining the disturbance inputs, as reported in Table A.7. The resulting upper bounds μ_{UB} are reported in Table 2.

Table 2: DFACS stability and performance upper bounds.

	Robust Stability	Robust Performance
μ_{UB}	0.442	0.967

Since $\mu_{UB} < 1$ in both cases, robust stability and robust performance are verified for the adopted linear model and prescribed uncertainty set, with the performance bound lying closer to unity and therefore indicating a more restrictive condition.

Delays are incorporated either deterministically through the nominal delay representation of the implemented discrete-time architecture or as uncertain dynamics using second-order Padé approximations. A noise breakdown is also performed using the linear time-invariant LISA closed-loop model to identify the dominant contributors to each performance channel.

The prescribed spectral masks define frequency-domain envelopes and do not require the underlying noise processes to be Gaussian. Within this framework, individual sources are propagated separately through the closed-loop system, while the total noise is reconstructed assuming mutually uncorrelated stochastic sources, consistently with the available mission noise baseline. When noise cross-correlations are available, they can instead be included through a full cross-spectral density matrix $S_w(\omega)$. For a closed-loop transfer matrix $T_{yw}(s)$,

$$S_y(\omega) = T_{yw}(j\omega) S_w(\omega) T_{yw}^*(j\omega), \quad (11)$$

thereby accounting for non-zero cross-spectral terms. Physical inter-channel couplings are addressed separately in the MIMO robustness analysis through sensing and actuation cross-talks and the associated structured uncertainty blocks.

Higher-order features of strongly non-Gaussian disturbances, such as impulsive or heavy-tailed events, cannot be characterized by spectral densities

alone and therefore require complementary nonlinear or statistically driven time-domain simulations when suitable stochastic models are available.

Within this spectral analysis, Fig. 10 shows the individual noise contributions to the TM position jitter together with the corresponding requirement.




Figure 10: Noise contributions to the TM linear position.

An additional robust-performance analysis evaluates the effect of model uncertainties on the TM position relative to the cage center. The performance channels are normalized with respect to their requirements, and Fig. 11 shows the corresponding singular values and uncertainty dispersion.




Figure 11: TM linear position singular values.

The scatter plots identify the frequency ranges most affected by uncertainties, while the noise breakdown highlights the dominant performance

contributors, supporting possible MIMO-level refinements.

Finally, high-fidelity MIMO validation is carried out through an extended Simulink simulation under worst-case input disturbance scenarios. The simulation spans 5×10^5 s, approximately 5.79 days, with the corresponding results reported in Fig. 12, where the black curve represents the y - and z -axis TM-to-optical-bench translational-jitter requirement reported in Table A.7.

The low-frequency violation on the y axis is primarily attributed to the high GRS sensing noise combined with the limited authority of the associated actuators, and this limitation is expected to be reassessed as updated GRS sensing-noise characterizations become available from ESA.




Figure 12: Amplitude Spectral Densities of TM to optical bench translational jitter.

5. Case Study II: The NGGM Mission

The Next Generation Gravity Mission (NGGM) is an ESA gravimetric mission concept proposed as a follow-up to GOCE, GRACE, and GRACE Follow-On, expected to be launched not earlier than 2032 [33] [34]. Its main objective is to monitor temporal variations of the Earth's gravity field with improved spatial and temporal resolution compared with previous satellite gravimetry missions [35], with these variations mainly driven by mass redistribution processes including hydrological cycles, ocean mass variability, ice-sheet and glacier mass loss, and solid-Earth deformation. For this reason, NGGM is meant not only to improve the long-term monitoring of climate-related mass transport phenomena, but also to enhance the observation of faster gravity-field variations occurring over a few days [36] [37] [38].

The mission concept relies on precise inter-satellite laser interferometry and ultra-sensitive accelerometers to separate the gravitational signal from

non-gravitational accelerations acting on the spacecraft [39] [34]. Accurate gravity-field reconstruction therefore requires both precise ranging and a sufficiently quiet dynamical environment for the accelerometer measurements, with non-gravitational disturbances sufficiently attenuated within the mission measurement bandwidth (MBW), defined here as 1–100 mHz [34]. This motivates the adoption of a DFACS based on micro-propulsion actuation and high-accuracy sensing [34] [40].

The complete NGGM control architecture includes coupled functions related to drag-free performance, angular stability, and pointing accuracy [34]; however, to assess the proposed methodology in a simplified and clearly interpretable setting, this case study is restricted to the Linear Drag-Free loop [40].

5.1. NGGM science-mode overview

In science mode, NGGM retrieves information on the Earth's gravity field by observing the relative motion of two spacecraft flying in a near-polar, low-Earth orbit and arranged in an in-line configuration [33]. The measurement principle, schematically shown in Fig. 13, is based on the Low-Low Satellite-to-Satellite Tracking (LL-SST) technique [35] [34].




Figure 13: Low-Low Satellite-to-Satellite Tracking (LL-SST) measurement principle for the NGGM mission [35].

As the two spacecraft occupy slightly different positions along the same orbital path, they are subject to different gravitational accelerations, which produce small variations in the inter-satellite distance measured with high precision by laser interferometry [35]. Gravity-field information is therefore

reconstructed from the relative distance variations between the two spacecraft. However, the inter-satellite ranging observable is also affected by non-gravitational perturbations, with atmospheric drag representing the dominant disturbance in low-Earth orbit, while solar radiation pressure, thermoelectric effects, and actuator-induced disturbances generally play a secondary role [34]. If not adequately compensated, these perturbations may introduce spurious signatures in the frequency band relevant for gravity-field recovery. Science-mode performance therefore requires a sufficiently quiet dynamical environment for the accelerometer measurements, together with the pointing stability needed to preserve the laser-link alignment [33].

This requirement is addressed by a DFACS architecture integrating payload sensors, spacecraft sensors, and micro-propulsion actuators, in which the accelerometers measure the residual non-gravitational accelerations acting on the spacecraft, while the control system commands the MPS to generate the forces and torques required to reduce their impact on the science measurement [34] [35].

The controller must therefore provide the required performance while remaining robust to environmental disturbances, sensor and actuator noise, and actuator limitations; accordingly, the science-mode architecture is organized into three coordinated control functions:

1. **Linear Drag-Free Loop.** The linear drag-free loop acts on the translational dynamics of the spacecraft center of mass. The accelerometer measures the residual non-gravitational accelerations, while the micro-propulsion system provides the compensating forces. Its objective is to minimize the residual translational accelerations within the MBW, thereby preserving the quality of the gravity-field observable.
2. **Angular Drag-Free Loop.** The angular drag-free loop controls the rotational dynamics of the spacecraft, reducing residual angular accelerations and rate disturbances that may couple into the accelerometer measurements. It therefore limits the interaction between translational and rotational dynamics and their impact on the reconstructed non-gravitational acceleration.
3. **Attitude and Pointing Control.** The attitude and pointing control loop maintains the spacecraft orientation required by the science payload and the inter-satellite laser link. Accurate pointing preserves the line-of-sight (LoS) alignment between the spacecraft and limits angular misalignments affecting the ranging observable.

As in the LISA case, these functions are coupled elements of the same science-mode architecture rather than independent control problems. Within the MBW, the main objective is to reduce residual linear and angular non-gravitational accelerations affecting the gravity-field measurement chain, whereas outside this bandwidth the requirements are generally relaxed to avoid unnecessary actuator activity and amplification of accelerometer and thruster noise. The controller must therefore provide sufficient disturbance attenuation within the scientific frequency range while preserving robustness to sensor noise, actuator limitations, and unmodeled dynamics.

5.2. Design results

To evaluate the effectiveness of the methodology introduced in Section 3, the NGGM linear drag-free loop is designed and assessed in science mode.

Remark 4. *At this stage, the analysis is restricted to the linear drag-free loop, so that the controller can be evaluated independently from the other control functions and the main frequency-domain trade-offs can be clearly identified. Within this framework, the objective is to reduce the residual non-gravitational acceleration within the MBW while limiting measurement-noise amplification and excessive micro-thruster activity.*

The control design follows the decoupled SISO scheme introduced in Fig. 5. After neglecting cross-axis couplings only at the preliminary synthesis stage, the translational dynamics is represented by three independent channels, one for each controlled axis, while the neglected coupling effects are subsequently represented in the MIMO robustness model through structured sensing and actuation cross-coupling uncertainties.

The frequency-domain requirements adopted for the NGGM linear drag-free design are summarized in Table B.9 (see the Appendix). Since this case study is limited to the linear drag-free loop, only the residual linear acceleration, thruster force noise, and accelerometer measurement-noise specifications are considered. The residual-acceleration ASD defines the main performance requirement within the MBW, while the thruster and accelerometer noise masks represent the main actuation and sensing limitations. In particular, thruster noise acts as an actuator-side disturbance, whereas accelerometer noise enters through the measurement chain and may be amplified by an overly aggressive controller. The resulting design therefore involves a frequency-domain trade-off between disturbance attenuation, sensor-noise limitation, and control effort.

Accordingly, the controller is synthesized by combining soft and hard constraints following the procedure introduced in Section 3. Soft constraints shape the main closed-loop propagation channels by limiting the effects of sensor noise and actuator-side disturbances on the residual acceleration, regulation error, and control input.

Hard constraints are instead assigned to operationally critical channels. In particular, the sensor-noise-to-control-input channel is constrained to limit the conversion of measurement noise into excessive micro-thruster commands, thereby contributing to the required in-band drag-free performance while limiting control activity associated with sensing noise. Actual actuator magnitude and rate limitations are assessed separately during the high-fidelity time-domain verification.

As shown in Fig. 14a, the first soft requirement is imposed on the sensor-noise-to-residual-acceleration transfer function, $n_{\text{sen}} \rightarrow y$, where n_{sen} denotes the accelerometer measurement noise and y the residual non-gravitational acceleration. This constraint limits the propagation of measurement noise into the drag-free performance output.

The second soft requirement, reported in Fig. 14b, is imposed on the sensor-noise-to-error transfer function, $n_{\text{sen}} \rightarrow e$. Since the error signal e is used by the controller to compute the corrective action, this constraint limits the propagation of spurious sensor fluctuations into the control action, avoiding unnecessary actuator activity and noise amplification, especially outside the measurement bandwidth.




Figure 14: Soft constraints defined for the drag-free loop.

The third soft requirement, reported in Fig. 15a, is imposed on the actuator-disturbance-to-output transfer function, $d_{\text{act}} \rightarrow y$. This channel quantifies the effect of actuation-side disturbances on the residual non-gravitational acceleration and directly shapes the disturbance-rejection capability of the controller within the MBW.

Finally, the fourth soft requirement, shown in Fig. 15b, is imposed on the actuator-disturbance-to-control-input transfer function, $d_{\text{act}} \rightarrow u$. This constraint limits the control effort generated in response to actuation-side disturbances and keeps it compatible with the available actuator authority. Together with $d_{\text{act}} \rightarrow y$, it shapes the trade-off between disturbance rejection and control activity.




Figure 15: Additional soft constraints defined for the drag-free loop.

Overall, the soft constraints shape the nominal closed-loop response by limiting the main noise and disturbance propagation paths, while the frequency-domain results show that all corresponding closed-loop gains remain below their prescribed bounds, indicating that the nominal design achieves the desired trade-off between disturbance attenuation, noise limitation, and control activity.

A hard constraint is imposed on the sensor-noise-to-control-input transfer function, $n_{\text{sen}} \rightarrow u$, which quantifies how measurement noise is converted into micro-thruster commands. This limitation is treated as non-negotiable because excessive actuation driven by sensing noise could exceed the available actuator authority. The constraint therefore prevents the controller from achieving drag-free performance at the cost of excessive thrust activity.




Figure 16: Hard constraint for the drag-free loop.

As shown in Fig. 16, the closed-loop gain remains below the prescribed bound over the considered frequency range.

5.3. Robustness verification results

The parametric uncertainties considered in the μ -analysis are summarized in Table B.10 of the Appendix and, consistently with the LFT-based robustness framework introduced in Section 3.2, are grouped into three physical families: spacecraft mass, micro-propulsion actuation, and accelerometer sensing. Note that, such values are provided by ESA according to the current development and state-of-the-work of the NGGM mission. Such values are likely to change in the future, depending on advancements and

The spacecraft mass uncertainty is introduced through the inverse-mass term $1/m$, which directly affects the force-to-acceleration conversion; since m is a scalar parameter, its uncertainty is represented by a repeated real scalar block acting identically along the three translational axes, without introducing independent axis-wise variations or cross-axis couplings.

Conversely, the MPS and accelerometer uncertainty blocks are modeled as 3×3 matrices. Their diagonal entries represent independent scale-factor errors along the controlled axes, described by the uncertain parameters p_i , while the off-diagonal terms p_{ij} , $i \neq j$, account for misalignments and cross-coupling effects between axes. These terms are expressed through the uncertain coefficients α_{ij} and scaled by the corresponding maximum cross-coupling amplitudes θ_{MPS} and θ_{Acc} .

This uncertainty representation is consistent with the decoupled synthesis adopted in the previous section, where the nominal controller is designed using the dominant diagonal translational dynamics, while the neglected cross-axis coupling effects are represented during MIMO robustness verification through structured sensing and actuation uncertainties. Thus, cross-axis couplings are treated as bounded deviations from the nominal decoupled configuration rather than assumed to be identically zero.

The adopted 3×3 uncertainty blocks represent sensing and actuation misalignments, cross-talk effects, and scale-factor variations within the prescribed bounds, so that the resulting robustness guarantees apply to coupling effects represented within this structured uncertainty set and over the analyzed frequency range. More complex frequency-dependent linear couplings can be included through dynamic uncertainty blocks or explicit MIMO channels, while nonlinear coupling effects are addressed through the subsequent high-fidelity time-domain verification.

As a preliminary robustness assessment, the nominal gain and phase margins of the representative SISO channel are evaluated from its open-loop transfer function. Although they do not account for the full structured uncertainty set, they provide an indication of nominal robustness, with the corresponding results reported in Table 3.

Table 3: Nominal DFACS stability margins.

Loop	Phase Margin [°]	Gain Margin [dB]
Linear DF	53.8	10.7

Subsequently, a structured singular value μ -analysis is performed to assess robust stability and robust performance under the modeled structured uncertainty set.

To ensure consistency with the formulation subsequently adopted in the μ -synthesis, robust performance is evaluated by augmenting the uncertainty structure with a single fictitious full-complex performance block associated with the complete normalized map from the exogenous inputs to the performance outputs. The resulting peak upper bound on μ are reported in Table 4.

Table 4: DFACS robust stability and robust performance μ upper bounds.

	Robust Stability	Robust Performance
μ_{UB}	0.5680	0.9764

Since the peak upper bounds remain below unity, the implemented drag-free loop satisfies both robust stability and robust performance for the considered uncertainty set, although the robust-performance upper bound remains close to the unit threshold.

To further interpret the frequency-domain results, a noise breakdown is performed on the linear time-invariant closed-loop model to identify the dominant disturbance and noise contributions to the residual linear acceleration, with particular attention to accelerometer noise and actuator-side disturbances, which directly affect the drag-free performance variable within and around the MBW. Fig. 17 reports the individual contributions together with the reconstructed total response. As in the LISA case, the individual sources




Figure 17: Noise contribution to the DF loop.

are propagated separately, while the total response is reconstructed assuming mutually uncorrelated disturbance and sensing-noise sources.

In addition, a worst-case performance analysis evaluates the sensitivity of the implemented drag-free loop to the modeled structured uncertainties. The relevant closed-loop performance channels are directly compared with the prescribed drag-free requirements, identifying the frequency regions most affected by admissible parameter variations.




Figure 18: Worst-case performance analysis for DF loop.

The corresponding results are reported in Fig. 18, comparing the nominal response with the responses obtained under the worst-case uncertainty realizations identified by the analysis. This analysis verifies whether the residual linear-acceleration requirement remains satisfied under the identified worst-case conditions.

5.4. μ -synthesis approach

The baseline $\mathcal{H}_\infty$ controller designed in Section 5.2 already satisfies the nominal drag-free requirements and provides the desired frequency-domain behavior within the mission measurement bandwidth.

Moreover, Table 4 confirms robust stability and robust performance for the considered structured uncertainty set. However, the robust-performance upper bound, $\mu_{\text{UB}}^{\text{RP}} = 0.9764$, is relatively close to unity, indicating a limited robustness margin.

Remark 5. *The use of μ -synthesis should not be interpreted as a replacement for the baseline $\mathcal{H}_\infty$ design, which already represents a feasible baseline solution satisfying the prescribed drag-free requirements. Instead, μ -synthesis is introduced as a robust refinement step to reduce sensitivity to structured uncertainties and improve worst-case robustness margins. The D–K iteration therefore aims to improve robustness while preserving compliance with the prescribed in-band requirements.*

The refinement is performed through fixed-structure μ -synthesis using the same uncertain LFT model and robust-performance formulation adopted in the preceding verification. The controller structure and order are fixed a priori and are not increased by the D–K iterations. The baseline $\mathcal{H}_\infty$ controller initializes the prescribed tunable controller structure, providing

$$\mu_{\text{UB,init}}^{\text{RP}} = 0.9756,$$

essentially identical to the baseline value $\mu_{\text{UB}}^{\text{RP}} = 0.9764$. The subsequent D–K iterations aim to reduce the peak μ of the augmented robust-performance problem while preserving the nominal drag-free requirements.

Table 5 reports the D–K iteration process. The value 1.1960 at iteration 1 does not represent the robust performance of the baseline or initialized controller, but the peak μ obtained during the D-step for the controller resulting from the first K-step.

Table 5: D-K iteration process for μ -synthesis.

Iteration	K Step	Peak μ	D Fit	D-Scaling Fit Order
1	6.809×10^7	1.1960	1.2050	144
2	0.6175	0.5774	0.5786	180
3	0.5658	0.5636	0.5668	204
4	0.5656	0.5624	0.5658	192
5	0.5658	0.5638	0.5671	216

The largest improvement occurs between iterations 1 and 2, where the peak μ decreases from 1.1960 to 0.5774. Subsequent iterations provide only minor improvements, with the minimum value

$$\mu_{\text{peak}} = 0.5624$$

obtained at iteration 4. At iteration 5, the peak slightly increases to 0.5638 while the D-fit order increases from 192 to 216. According to the termination criterion introduced in Section 3, further refinement is not considered beneficial, and iteration 4 is selected.

The D-scaling fit orders reported in Table 5 do not correspond to the implemented controller order. The values from 144 to 216 refer only to rational approximations of the frequency-dependent D -scalings used internally during the offline D-K iteration and are not part of the on-board controller. Since the K-step is restricted to the fixed parameterization in Eq. (7), the implemented controller remains fourth order for each controlled SISO channel. Thus, increasing D-fit order affects only the offline synthesis complexity, not the real-time computational burden.

Table 6 summarizes the effect of the μ -synthesis refinement.

Table 6: Comparison of μ_{UB} before and after μ -synthesis.

	Robust Stability	Robust Performance
μ_{UB} after $\mathcal{H}_\infty$	0.5680	0.9764
μ_{UB} after μ -synthesis	0.3301	0.5613

The baseline and refined controllers are evaluated using the same robust-stability and robust-performance formulations. The selected controller provides $\mu_{\text{UB}}^{\text{RP}} = 0.5613$, in close agreement with the peak $\mu = 0.5624$ obtained

during the D–K iteration. The robust-stability upper bound decreases from 0.5680 to 0.3301, while the robust-performance upper bound decreases from 0.9764 to 0.5613, remaining below unity in both cases.

Overall, μ -synthesis reduces the robust-performance upper bound by approximately 42.5% while preserving the prescribed in-band requirements and the low-order controller structure.

5.5. Spectral performance verification

To complement the verification of the implemented controller, the nominal closed-loop response is assessed through the amplitude spectral density (ASD) of the residual linear acceleration. The simulation spans 250 orbits, approximately 1.3719×10^6 s (15.9 days), and refers only to the implemented linear drag-free loop.




Figure 19: Amplitude Spectral Density (ASD) of the residual linear acceleration.

The residual acceleration ASD is compared with the NGGM linear drag-free frequency-domain requirement reported in Table B.9 and shown in Fig. 19, remaining below the prescribed bound throughout the MBW and confirming in-band compliance for the considered linear loop.

Spectral components outside the MBW, particularly low-frequency narrow-band tone-error peaks, are also relevant and are typically associated with the orbital frequency and its harmonics [41]. When propagated from residual acceleration to relative displacement through double integration, their effect is amplified at low frequencies and may generate significant long-period displacement oscillations.

These effects are only partially addressed here because the implemented architecture is restricted to the linear drag-free loop. A complete mission-level architecture should also include higher-level functions for long-term inter-satellite formation regulation. Therefore, although the in-band requirement is satisfied, the out-of-band tone-error content motivates complementary spectral constraints and control functions in future developments.

6. Conclusion

This paper presents a systematic approach to robust control design and verification for critical drag-free space missions. The proposed methodology is based on a multi-stage Drag-Free Attitude Control System (DFACS) design process, combining SISO control synthesis, MIMO-level refinement when needed, and robustness, stability, and performance verification.

The main contribution lies in the integration of established robust-control tools within a hierarchical and iterative workflow, rather than in the introduction of a new individual control algorithm. In particular, controller complexity is progressively increased only when required by the verification process.

The control design relies on an $\mathcal{H}_\infty$ framework, suitable for addressing demanding scientific and system-level requirements in the presence of disturbances, sensor noise, control-effort constraints, and model uncertainties. Robust stability and robust performance are then verified through μ -analysis with respect to the prescribed structured uncertainty sets, providing information to identify critical channels and determine whether further refinement through μ -synthesis is required.

The methodology has been illustrated through two representative drag-free mission scenarios, LISA and NGGM, characterized by different control architectures, performance objectives, and uncertainty descriptions. These case studies demonstrate how the same hierarchical design and verification logic can be adapted to mission-specific requirements while retaining the overall workflow.

Overall, the proposed framework provides a structured procedure for the design and verification of high-precision drag-free control systems. By preserving the same overall sequence of SISO synthesis, MIMO robustness assessment, optional robust refinement, and high-fidelity verification, while adapting the plant model, mission requirements, disturbance descriptions,

and uncertainty sets, the methodology can be extended to a broad range of drag-free mission scenarios.

Declaration of competing interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

- [1] E. Canuto, Drag-free and attitude control for the GOCE satellite, *Automatica* 44 (7) (2008) 1766–1780.
- [2] W. Zheng, H. Hsu, M. Zhong, al., Requirements analysis for future satellite gravity mission improved-GRACE, *Surv. Geophys.* 36 (2015) 87–109.
- [3] M. Boggio, P. Cotugno, C. Perez Montenegro, M. Pagone, C. Novara, L. Massotti, NMPC-based orbit and formation control for an earth-gravity monitoring mission, in: *Proc. 72nd Int. Astronaut. Congr. IAC*, 2021.
- [4] K. Danzmann, al., LISA: laser interferometer space antenna for gravitational wave measurements, *Classical Quant. Grav.* 13 (11A) (1996).
- [5] D. Shaddock, Space-based gravitational wave detection with LISA, *Classical Quant. Grav.* 25 (11) (2008).
- [6] S. Vidano, C. Novara, L. Colangelo, J. Grzymisch, The LISA DFACS: A nonlinear model for the spacecraft dynamics, *Aerosp. Sci. Technol.* 107 (2020) 106313.
- [7] M. Galullo, G. Bucchioni, G. Franzini, M. Innocenti, Closed loop guidance during close range rendezvous in a three body problem, *J. Astronaut. Sci.* 69 (2022) 28–50.
- [8] M. Leomanni, A. Garulli, A. Giannitrapani, F. Scortecci, All-electric spacecraft precision pointing using model predictive control, *J. Guid. Control Dyn.* 38 (1) (2015) 161–168.

- [9] S. Vidano, C. Novara, M. Pagone, J. Grzymisch, The LISA DFACS: Model predictive control design for the test mass release phase, *Acta Astronaut.* 193 (2022) 731–743.
- [10] S. Vidano, M. Pagone, J. Grzymisch, V. Preda, C. Novara, Drag-free and attitude control system for the LISA space mission: An H_∞ constrained decoupling approach, *IEEE Trans. Control Syst. Technol.* 32 (6) (2024) 2149–2163.
- [11] P. Touboul, B. Foulon, M. Rodrigues, J. Marque, In orbit nano-g measurements, lessons for future space missions, *Aerosp. Sci. Technol.* 8 (5) (2004) 431–441.
- [12] A. Peidou, S. Pagiatakis, Gravity gradiometry with GRACE space missions: New opportunities for the geosciences, *J. Geophys. Res. Solid Earth* 124 (8) (2019) 9130–9147.
- [13] E. Canuto, Embedded model control: Outline of the theory, *ISA Trans.* 46 (3) (2007) 363–377.
- [14] E. Canuto, C. Novara, L. Massotti, D. Carlucci, C. Perez Montenegro, *Spacecraft Dynamics and Control. The Embedded Model Control Approach*, Elsevier Aerospace Engineering Series, 2018.
- [15] L. Pettazzi, A. Lanzon, S. Theil, Design of robust decentralized controllers for drag-free satellite, in: *Proc. 47th IEEE Conf. Decis. Control*, 2008, pp. 602–607.
- [16] L. Pettazzi, A. Lanzon, S. Theil, A. E. Finzi, Design of robust drag-free controllers with given structure, *J. Guid. Control Dyn.* 32 (5) (2009) 1609–1620.
- [17] A. Falcoz, C. Pittet, S. Bennani, A. Guignard, C. Bayart, B. Frapard, Systematic design methods of robust and structured controllers for satellites: Application to the refinement of Rosetta’s orbit controller, *CEAS Space J.* 7 (3) (2015) 319–334.
- [18] Z. Ma, J. Wang, A controller design method for drag-free spacecraft multiple loops with frequency domain constraints, *IEEE Trans. Aerosp. Electron. Syst.* 59 (3) (2023) 3224–3235.

- [19] X. Sun, Q. Shen, S. Wu, Partial state feedback MRAC-based reconfigurable fault-tolerant control of drag-free satellite with bounded estimation error, *IEEE Trans. Aerosp. Electron. Syst.* (2023).
- [20] M. Pagone, C. Novara, P. Martella, C. Nocerino, GNC robustness stability verification for an autonomous lander, *Aerosp. Sci. Technol.* 100 (2020) 105831.
- [21] C. Novara, S. Vidano, M. Pagone, J. Grzymisch, V. Preda, LISA drag-free attitude control system: Robust stability and performance analysis, *Acta Astronaut.* 240 (2026) 115–126.
- [22] C. Novara, M. Virdis, M. Pagone, D. Ruggiero, E. Capello, E. Punta, S. Dionisio, S. Vidano, J. Grzymisch, V. Preda, Recovery strategies to cope with micrometeoroid impacts in the LISA mission, *Acta Astronaut.* 211 (2023) 844–864.
- [23] K. Zhou, J. C. Doyle, *Essentials of Robust Control*, Prentice Hall, 1998.
- [24] A. Packard, J. C. Doyle, The complex structured singular value, *Automatica* 29 (1) (1993) 71–109.
- [25] D. Navarro-Tapia, A. Marcos, Legacy-based design of the LISA mission science control mode, in: *Proc. IEEE Conf. Control Technol. Appl.*, 2023.
- [26] ECSS-E-HB-60-10A - Control performance guidelines, European Space Agency, 2010.
URL <https://ecss.nl/hbstms/ecss-e-hb-60-10a-control-performance-guidelines/>
- [27] K. Zhou, J. C. Doyle, *Essentials of Robust Control*, Prentice Hall, 1998.
- [28] S. Tøffner-Clausen, P. Andersen, Mu-synthesis: A non-conservative methodology for design of controllers with robustness towards dynamic and parametric uncertainty, in: *Recent Results in Robust and Adaptive Control*, EURACO Workshop Florence 11-14 September 1995, 1995, pp. 269–303, mu-Synthesis ; Conference date: 19-05-2010.
- [29] M. Colpi, al., LISA definition study report, assessment Study Report, ESA, Paris, France (2024). [arXiv:2402.07571](https://arxiv.org/abs/2402.07571).
URL <https://arxiv.org/abs/2402.07571>

- [30] K. Danzmann, al., LISA: Unveiling a hidden universe, eSA/SRE(2011)3 (2011).
URL https://sci.esa.int/documents/34985/36280/1567258287202-NGO_YB.pdf
- [31] B. Abbott, al., Observation of gravitational waves from a binary black hole merger, *Phys. Rev. Lett.* 116 (02 2016).
- [32] Z. Chen, Detection of gravitational wave in LIGO and significance evaluation, *Highl. Sci. Eng. Technol.* 38 (2023) 264–269.
- [33] R. Haagmans, C. Siemes, L. Massotti, O. Carraz, P. Silvestrin, ESA’s next-generation gravity mission concepts, *Rendiconti Lincei. Scienze Fisiche e Naturali* 31 (2020) S15–S25.
- [34] S. Cesare, S. Dionisio, M. Saponara, D. Bravo-Berguño, L. Massotti, J. Teixeira da Encarnação, B. Christophe, Drag and attitude control for the next generation gravity mission, *Remote Sens.* 14 (12) (2022) 2916.
- [35] L. Massotti, D. Di Cara, J. G. del Amo, R. Haagmans, M. Jost, C. Siemes, P. Silvestrin, The ESA earth observation programmes activities for the preparation of the next generation gravity mission, in: *AIAA GNC Conference, American Institute of Aeronautics and Astronautics, Boston, MA, USA, 2013.*
- [36] L. Colangelo, E. Canuto, L. Massotti, C. Novara, M. A. Lotufo, Orbit and formation control for the next generation gravity mission, *IFAC-PapersOnLine* 49 (17) (2016) 284–289.
- [37] L. Massotti, C. Siemes, G. March, R. Haagmans, P. Silvestrin, Next generation gravity mission elements of the mass change and geosciences international constellation: From orbit selection to instrument and mission design, *Remote Sens.* 13 (19) (2021) 3935.
- [38] M. Boggio, L. Colangelo, M. Viridis, M. Pagone, C. Novara, Earth gravity in-orbit sensing: MPC formation control based on a novel constellation model, *Remote Sens.* 14 (12) (2022).
- [39] K. Nicklaus, K. Voss, A. Feiri, M. Kaufer, C. Dahl, M. Herding, B. A. Curzadd, A. Baatzsch, J. Flock, M. Weller, V. Müller, G. Heinzel,

- M. Misfeldt, J. J. E. Delgado, Towards NGGM: Laser tracking instrument for the next generation of gravity missions, *Remote Sens.* 14 (16) (2022) 4089.
- [40] A. Bacchetta, L. Colangelo, E. Canuto, S. Dionisio, L. Massotti, C. Novara, M. Parisch, P. Silvestrin, From GOCE to NGGM: Automatic control breakthroughs for european future gravity missions, *IFAC-PapersOnLine* 50 (1) (2017) 6428–6433.
- [41] N. Pfaffenzeller, R. Pail, T. Gruber, Impact of tone errors in future satellite gravimetry missions, *Adv. Space Res.* 75 (10) (2025) 6961–6980.

Appendix A. LISA Performance Requirements

Table A.7: LISA signals specifications in the frequency band $0.1 \text{ mHz} \leq f \leq 0.1 \text{ Hz}$.

Signal		Frequency domain specification
<i>Channel Requirements</i>		
TM	y, z	$4.7 \sqrt{1 + \frac{10^{-3}}{f}} \frac{\sqrt{1 + \left(\frac{f}{8 \times 10^{-3}}\right)^4}}{\sqrt{1 + \left(\frac{f}{8 \times 10^{-3}/\sqrt{2}}\right)^4}} \text{ nm Hz}^{-1/2}$
	η, ϕ	$10 \sqrt{1 + \left(\frac{4 \times 10^{-3}}{f}\right)^2 + \left(\frac{0.5 \times 10^{-3}}{f}\right)^2} \text{ nrad Hz}^{-1/2}$
SC	θ, η	$8 \sqrt{1 + \left(\frac{1.2 \times 10^{-3}}{f}\right)^2} \text{ nrad Hz}^{-1/2}$
	ϕ	$4 \sqrt{1 + \left(\frac{1.2 \times 10^{-3}}{f}\right)^2} \text{ nrad Hz}^{-1/2}$
MOSA	ϕ	$4 \sqrt{1 + \left(\frac{1.2 \times 10^{-3}}{f}\right)^2} \text{ nrad Hz}^{-1/2}$
GRS	u_y	$0.2 \times 10^{-12} \sqrt{1 + \left(\frac{5 \times 10^{-5}}{f}\right)^2} \text{ m s}^{-2} \text{ Hz}^{-1/2}$
	u_z	$0.1 \times 10^{-12} \sqrt{1 + \left(\frac{7 \times 10^{-5}}{f}\right)^2} \text{ m s}^{-2} \text{ Hz}^{-1/2}$
<i>Noise and Disturbance Specifications</i>		
GRS Actuation	y, z	$4.5 \times 10^{-15} \sqrt{\left(\left(\frac{10^{-4}}{f}\right)^2 + 0.175^2\right) \left(1 + \left(\frac{f}{8 \times 10^{-3}}\right)^4\right)} \text{ m s}^{-2} \text{ Hz}^{-1/2}$
GRS Sensor	y	$1.8 \times 10^{-9} \sqrt{1 + \left(\frac{10^{-3}}{f}\right)^2} \text{ m Hz}^{-1/2}$
	z	$3 \times 10^{-9} \sqrt{1 + \left(\frac{10^{-3}}{f}\right)^2} \text{ m Hz}^{-1/2}$

Table A.8: LISA parametric uncertainties considered in the μ -analysis. Each parameter is modeled as $p = p_0(1 + \delta)$, where p_0 denotes the nominal value, p the uncertain parameter, and δ a real scalar uncertainty satisfying $|\delta| \leq \bar{\delta}$.

Uncertain source	Uncertainty model
TM mass	$p = p_0(1 + \delta), \delta \leq 0.01$
TM inertia	$p = p_0(1 + \delta), \delta \leq 0.01$
SC mass	$p = p_0(1 + \delta), \delta \leq 0.10$
SC inertia	$p = p_0(1 + \delta), \delta \leq 0.10$
TM stiffness	$p \in [0, 1.05] p_0$
GRS actuation crosstalk	$p \in [0, 1.05] p_0$
GRS sensing crosstalk	$p \in [0, 1.05] p_0$
GRS actuation gain	$p \in [0.8, 1.2] p_0$
GRS sensing gain	$p \in [0.8, 1.2] p_0$
MPS actuation gain	$p = p_0(1 + \delta), \delta \leq 0.10$

Appendix B. NGGM Performance Requirements

Table B.9: NGGM linear drag-free frequency domain specifications and spectral masks used for the control design and verification. The mission measurement bandwidth is defined as $f_1 = 10^{-3}$ Hz and $f_2 = 10^{-1}$ Hz, i.e., 1–100 mHz.

Function / Signal	Frequency domain specification	Parameters
Linear Drag-Free CoM residual acceleration ASD	$\text{ASD}_{a,\text{req}}(f) = A_0, \quad f_1 \leq f \leq f_2$	$A_0 = 5 \times 10^{-9} \text{ m s}^{-2} \text{ Hz}^{-1/2}$
Thruster noise Thruster force-noise ASD	$\text{ASD}_F(f) = \begin{cases} d_0, & f \leq f_L, \\ d_0 \frac{f_L}{f}, & f_L < f < f_H, \\ d_H, & f \geq f_H, \end{cases}$	$d_0 = 30 \times 10^{-6} \text{ N Hz}^{-1/2}$ $f_L = 3 \text{ mHz}$ $f_H = 75 \text{ mHz}$ $d_H = 1.2 \times 10^{-6} \text{ N Hz}^{-1/2}$
Accelerometer noise Accelerometer-noise ASD	$\text{ASD}_n(f) = \begin{cases} n_0 \left(\frac{f_1}{f}\right)^2, & f < f_1, \\ n_0, & f_1 \leq f \leq f_2, \\ n_0 \left(\frac{f}{f_2}\right)^2, & f > f_2, \end{cases}$	$n_0 = 10^{-11} \text{ m s}^{-2} \text{ Hz}^{-1/2}$

Table B.10: NGGM parametric uncertainties considered in the μ -analysis. For diagonal gain and mass parameters, δ denotes a real scalar uncertainty. For off-diagonal cross-coupling terms, p_{ij} denotes the coupling from axis j to axis i , with $i \neq j$; α_{ij} is a real uncertain coefficient satisfying $\alpha_{ij} \in [-1.05, 1.05]$, while $\theta_{\text{MPS}} = 0.5^\circ$ and $\theta_{\text{Acc}} = 0.2^\circ$ define the maximum nominal cross-coupling amplitudes for the MPS and accelerometer blocks, respectively.

Uncertain source	Uncertainty model
SC1 mass	$m = m_0(1 + \delta)$, $ \delta \leq 0.10$
SC2 mass	$m = m_0(1 + \delta)$, $ \delta \leq 0.10$
SC1 MPS actuation gain	$p = p_0(1 + \delta)$, $ \delta \leq 0.10$
SC2 MPS actuation gain	$p = p_0(1 + \delta)$, $ \delta \leq 0.10$
SC1 MPS actuation crosstalk	$p_{ij} = \alpha_{ij}\theta_{\text{MPS}}$, $\alpha_{ij} \in [-1.05, 1.05]$, $i \neq j$, $\theta_{\text{MPS}} = 0.5^\circ$
SC2 MPS actuation crosstalk	$p_{ij} = \alpha_{ij}\theta_{\text{MPS}}$, $\alpha_{ij} \in [-1.05, 1.05]$, $i \neq j$, $\theta_{\text{MPS}} = 0.5^\circ$
SC1 accelerometer sensing gain	$p = p_0(1 + \delta)$, $ \delta \leq 0.10$
SC2 accelerometer sensing gain	$p = p_0(1 + \delta)$, $ \delta \leq 0.10$
SC1 accelerometer sensing crosstalk	$p_{ij} = \alpha_{ij}\theta_{\text{Acc}}$, $\alpha_{ij} \in [-1.05, 1.05]$, $i \neq j$, $\theta_{\text{Acc}} = 0.2^\circ$
SC2 accelerometer sensing crosstalk	$p_{ij} = \alpha_{ij}\theta_{\text{Acc}}$, $\alpha_{ij} \in [-1.05, 1.05]$, $i \neq j$, $\theta_{\text{Acc}} = 0.2^\circ$