doi: 10.82124/CEAS-GNC-2026-097

EuroGNC Madrid, Spain May 5th-7th 2026

STARGATE: Towards a Robust Guidance Strategy for Autonomous Systems

Riccardo Minnozzi Research Assistant, Alma Mater Studiorum - Universitá di Bologna, Forlí, Italy. riccardo.minnozzi@unibo.it

Marco Sagliano Full Professor, Alma Mater Studiorum - Universitá di Bologna, Forlí, Italy. marco.sagliano@unibo.it

ABSTRACT

As autonomous spacecraft take on increasingly critical roles in future missions, the ability to operate safely under uncertainty becomes a key factor towards mission success. Existing trajectory design techniques primarily rely on deterministic formulations, delegating the management of unplanned disturbances to feedback control systems. While effective for nominal operations, this reactive paradigm limits performance and cannot guarantee safety in the presence of significant modeling errors or environmental disturbances without significant performance degradation. The European Research Council (ERC)-funded STochastic Algorithm for Robust Guidance Analysis and Trajectory Estimation (STARGATE) project aims to address this limitation by embedding uncertainty directly within the trajectory optimization process, shifting from reactive robustness to proactive resilience. Building upon and extending the Unscented Guidance approach, the framework seeks to capture arbitrary stochastic processes through efficient uncertainty representations. Focusing on scalable and accurate optimization methods, the proposed approach plans to leverage Optimal Control Theory and Consensus Optimization to guarantee the numerical tractability required for onboard applications. The proposed methodology will be validated on representative hardware through state-of-the-art testbenches, including rocket powered landing and rendezvous scenarios. By generating trajectories that are inherently more robust to uncertainties, STARGATE aims to enhance both safety and autonomy of future space missions.

Keywords: Stochastic Optimal Control, Robust Autonomous Systems, Consensus Optimization

Nomenclature

DoF Degrees of Freedom 8

EAGLE Environment for Autonomous GNC Landing Experiments 8
ERC European Research Council 1, 5, 10

NLP Nonlinear Programming 4, 5

STARGATE STochastic Algorithm for Robust Guidance Analysis and Trajectory Estimation 1, 2, 5–10

TEAMS Test Environment for Applications of Multiple Spacecraft 8

1 Introduction

Uncertainty is an inherent feature of space missions and directly impacts the probability of mission success. In several high-profile cases, unmodeled disturbances, such as hardware degradation or unexpected environmental conditions, have forced costly corrective actions. For instance, during the Rosetta mission, a thruster anomaly detected after several years of flight required an extensive 3-year recovery effort to prevent mission failure [1]. While such interventions are manageable in isolated missions, they are incompatible with the emerging operational landscape, characterized by large satellite constellations [2] and increasingly frequent launches of reusable vehicles [3]. In this context, autonomous systems have become a key enabler for scalable mission operations, allowing for faster decision-making and reducing operational overhead.

The limited computational resources available on space-qualified processors impose tight constraints on trajectory generation algorithms. As a result, current practice relies on nominal models to describe system dynamics during the planning phase. While such simplifications are often justified, they become inadequate when environmental disturbances or modeling errors are non-negligible. In operational architectures, these discrepancies are handled reactively by feedback controllers, which are expected to correct deviations from the nominal trajectory as schematically illustrated in Figure 1.

However, feedback alone cannot guarantee mission safety under all conditions. Launch operations, for example, are frequently delayed when wind or atmospheric disturbances exceed acceptable limits [4], highlighting the limitations of purely reactive strategies. In addition, treating uncertainty exclusively as a disturbance to be rejected prevents the full exploitation of the available design space, often leading to overly conservative designs and unnecessary performance degradation. This motivates the need for developing trajectory-planning frameworks that incorporate uncertainty directly into the decision-making process.

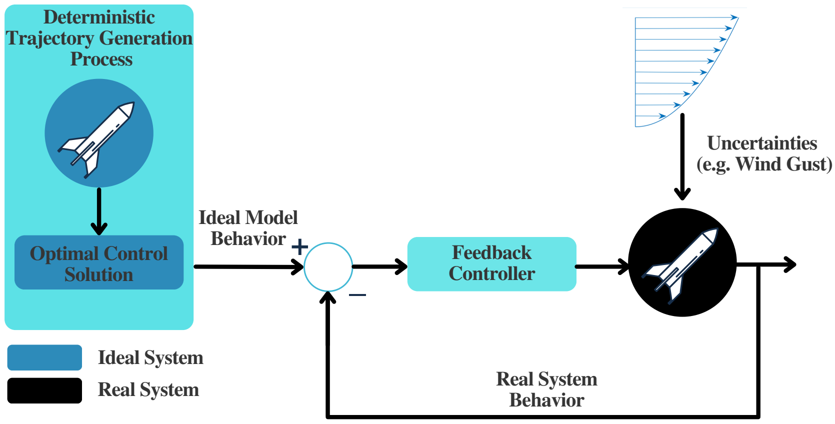

Fig. 1 Deterministic approach: the controller reacts to the uncertainties acting on the real system.

The current state of research in robust trajectory planning exhibits several key limitations, from the reliance on simplifying assumptions to the “curse of dimensionality”—the exponential growth in problem complexity with the number of decision variables. These challenges are particularly pronounced in realistic space applications, such as the landing of reusable launchers, where onboard computational resources are limited. In this paper we introduce the STARGATE framework, which aims at formulating efficient methodologies to embed uncertainties directly into the decision-making process and developing scalable, real-time-capable solution algorithms, ultimately demonstrating the approach on realistic hardware.

This paper is structured as follows. Section 2 surveys existing approaches to the problem, with Sec. 2.1 focusing on uncertainty modeling and its integration into trajectory optimization, and Sec. 2.2 discussing solution strategies for the resulting optimal control problems, emphasizing scalability and numerical tractability. Section 3 introduces the STARGATE framework, outlining its methodological foundations, their application to a representative benchmark test problem, and the experimental platforms that will be used for validation. Finally, section 4 summarizes the main contributions and discusses future developments.

2 Current Approaches to Robust Trajectory Design

Robust trajectory design focuses on generating trajectories that remain feasible and performant in the presence of uncertainties arising from system dynamics, environmental disturbances, or model inaccuracies. Unlike traditional formulations, where the optimal control problem is solved based on a deterministic model, robust approaches embed uncertainty directly into the problem definition [5]. The remainder of this section reviews the main developments in this field, beginning with the formulation of uncertain optimal control problems in subsection 2.1 and followed by solution strategies in subsection 2.2, highlighting the limitations that motivate the need for new methodologies.

2.1 Uncertainty Modeling

The inclusion of uncertainty in optimal control formulations has been studied extensively, with numerous techniques proposed to propagate disturbances, approximate probability distributions, or enforce robustness constraints [5–7]. However, most of these approaches rely on strong structural assumptions or do not scale well in the case of more realistic and challenging problems. Early efforts based on model augmentation, such as shadow trajectory methods [8], embed uncertainty into an indirectly-solved problem, but lead to prohibitively long solution times, even for low-dimensional problems. Covariance-steering formulations [9] reduce computational effort but remain constrained by poor scalability in higher-dimensional settings as well as in terms of the nature of the stochastic distributions that can be dealt with.

Other strategies aim to approximate the impact of uncertainty through sensitivity-based or distributional representations. Desensitized optimal control [10, 11] introduces sensitivity dynamics into the problem formulation, but the number of additional differential equations grows rapidly with the size of the uncertainty space. Gaussian Mixture Models [12–15] provide a flexible way to approximate non-Gaussian disturbances, but accurate representations require many mixture components, computationally intensive selection schemes [13, 14], and potentially multistage decision-making [15], thus complicating the computational and implementational aspects of the approach. Recent work on Koopman-based linearization [16] offers an alternative route for nonlinear systems, yet still suffers from dimensionality and representation issues in strongly nonlinear regimes. These limitations have motivated the development of more structured approaches to uncertainty modeling, notably Polynomial Chaos Theory and Unscented Guidance.

2.1.1 Polynomial Chaos Theory

In Polynomial Chaos Theory [17–21], the uncertain parameters are replaced by their expansion based on orthogonal polynomials. Instead of treating the system parameters as stochastic entities, each uncertain quantity is expressed as a finite series of deterministic coefficients applied to orthogonal polynomials of random variables. When this expansion is substituted into the system dynamics, the original stochastic optimal control problem is transformed into a larger, fully deterministic one, allowing the propagation of uncertainty without resorting to sampling.

Polynomial Chaos Theory has been successfully applied in several studies [20, 21], yet it presents key limitations that hinder its applicability to complex systems. First, the augmented optimal control problem requires explicit re-derivation whenever the number of uncertain parameters or the order of the expansion changes, limiting the flexibility of the method. Second, because the expansions are polynomial in nature, nonlinear dynamics often require reformulation or approximation to preserve tractability, which may compromise accuracy. Finally, Polynomial Chaos methods inherently suffer from the curse of dimensionality: the number of expansion terms grows combinatorially with the number of uncertain variables and the chosen polynomial order, rapidly making the resulting problem numerically intractable for real-time or high-dimensional applications [22]. Despite some attempts to mitigate such drawbacks [23], these challenges call for the exploration of alternative uncertainty-propagation strategies, such as Unscented Guidance (or σ-point-based) methods.

2.1.2 Unscented Guidance

Unscented Guidance [24–31] adopts a σ-point-based representation of uncertainty derived from the Unscented Transform [32]. The Unscented Transform approximates the effect of stochastic variables by deterministically selecting a small set of representative σ-points and their associated weights. For an n-dimensional Gaussian distribution, only $2n + 1$ points are required to capture the mean and covariance, providing an efficient alternative to sampling-based uncertainty propagation.

Originally developed for state estimation and navigation [33], the Unscented Transform was later extended to trajectory optimization, enabling the computation of control profiles robust to Gaussian process uncertainties [25–28]. Further generalizations have extended this principle to uniform (through the conjugate approach) [34, 35] and non-Gaussian [36, 37] distributions through modified σ-point selection schemes. Compared to Polynomial Chaos methods, Unscented Guidance scales more favorably with the number of uncertain parameters while maintaining accuracy in capturing low-order stochastic moments.

Despite its advantageous features, Unscented Guidance still presents key limitations that restrict its general applicability. The Conjugated Unscented Transform is limited to Gaussian and uniform distributions, and the number of σ-points required to accurately capture higher-dimensional or higher-order stochastic moments grows rapidly [34]. More recent generalizations mitigate these restrictions only partially, as they often rely on assumptions about the underlying probability density or cannot consistently represent higher-order moments [35, 36]. Furthermore, the most advanced formulations do not constrain the σ-weights to lie within the $[0, 1]$ interval [37]. Since these weights act as numerical quadrature coefficients in the multidimensional integration of stochastic moments, unconstrained values may lead to numerical ill-conditioning or loss of integration accuracy [38].

Finally, while Unscented Guidance provides advantageous scalability properties, it introduces complications from a numerical standpoint. The system dynamics must be propagated in parallel for all σ-points, while a single set of control inputs is shared among them [24–26]. This coupling introduces strong dependencies across the differential equations, reducing problem sparsity and complicating the formulation of a Nonlinear Programming (NLP) problem suitable for standard optimization solvers [39, 40], potentially leading to major problems when increasing the number of uncertainty sources [41, 42].

2.2 Solution Methods

The solution of the formulated robust optimal control problem aims to determine a control policy that optimizes a performance index while accounting for the stochastic evolution of the system dynamics. When applied to trajectory optimization, this task rapidly becomes computationally demanding, as each uncertainty source increases the number of trajectories to be propagated and the degree of coupling among them [24–26]. The resulting optimization problems are typically large, nonlinear, and highly constrained, requiring efficient numerical strategies to enable onboard applications.

The earliest class of solution techniques is represented by indirect methods. These approaches derive the necessary conditions for optimality (from Pontryagin's Minimum Principle [43]) in the form of a two-point boundary value problem. Their analytical rigor and high accuracy make them attractive for low-dimensional problems and for verifying the optimality of direct or approximate solutions [26, 44].

Indirect methods have been investigated in the context of robust trajectory design through stochastic extensions of Pontryagin's Minimum Principle, as well as in covariance-steering problems [9]. However, their practical adoption remains limited by sensitivity to initial guesses, numerical instability in the presence of nonlinearities, and poor scalability to high-dimensional stochastic systems [8].

In aerospace contexts, direct methods are widely adopted for trajectory optimization [45–47]. These techniques rely on a discretization of the system dynamics to generate a corresponding NLP, offering robustness and compatibility with complex dynamics and constraints. However, when uncertainties are embedded into the formulation, the dimensionality of the NLP grows significantly, causing an increase in memory requirements and the deterioration of numerical conditioning [48], thus negatively affecting convergence. Direct methods have been successfully applied to the solution of low-dimensional robust trajectory optimization problems [24].

More recently, consensus-based approaches have been proposed. These methods partition the global problem into smaller subproblems, each associated with a subset of decision variables, iteratively enforcing the "agreement" among the subproblems through coordination schemes [49]. The structure of consensus-based optimization allows parallelization [50], improving its scalability properties, at the cost of limited solution accuracy [51].

Overall, current solution techniques reveal a persistent trade-off between accuracy and scalability: indirect methods provide analytical rigor but lack robustness for complex problems, direct methods ensure reliability at the cost of computational efficiency, and consensus schemes enable distributed computation at the expense of solution accuracy. These limitations motivate the development of alternative approaches, capable of providing both scalability and accuracy.

3 STARGATE framework

The current state of robust trajectory design highlights limitations shared across the field, with approaches to problem formulation and solution being constrained by computational costs and structural rigidity. Techniques such as Polynomial Chaos expansions [17–21] and Unscented Guidance [24–31] have demonstrated the feasibility of embedding uncertainty into the trajectory optimization process, yet their scalability to realistic, high-dimensional problems remains limited. On the other hand, current numerical solvers struggle to efficiently handle the coupling and dimensionality introduced by stochastic models. As a result, operational systems continue to rely on deterministic trajectories, leveraging feedback control to address unplanned disturbances.

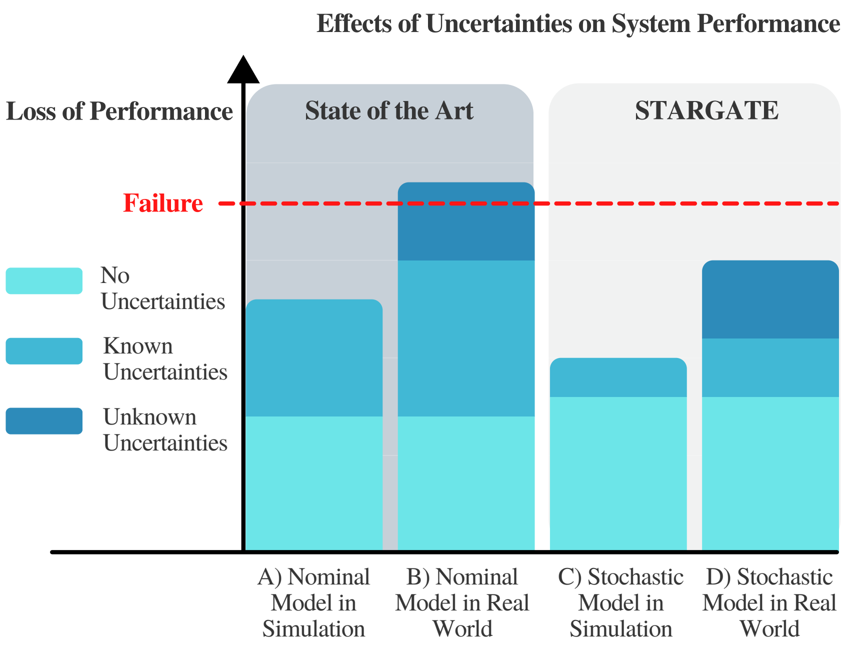

Fig. 2 Target behavior of STARGATE-enabled autonomous systems.

To address this gap, the ERC-funded STARGATE project investigates a new paradigm for robust trajectory planning: embedding uncertainty directly within the optimal control formulation to generate trajectories that are proactively resilient, rather than reactively corrected. A schematic illustration of the performance and safety benefits of uncertainty-aware trajectory planning is provided in Figure 2. The central research question driving this effort is therefore:

How can optimal control theory be expanded to systematically include uncertainties in the process and rapidly generate stochastically optimal solutions able to exploit them?

To realize the envisioned shift to a robust trajectory planning paradigm, STARGATE is structured around three complementary research objectives. The first, introduced in subsection 3.1, focuses on extending uncertainty modeling techniques to accurately represent different stochastic processes and efficiently embedding them within optimal control problems. The second, presented in subsection 3.2, introduces a multidisciplinary approach to the solution process of the resulting optimal control problem, prioritizing scalability and enhancing accuracy and convergence speed. Finally, the third objective, illustrated in subsection 3.3, is to validate the proposed framework on representative hardware, demonstrating the feasibility of real-time, onboard implementation. Together, these components form a cohesive framework that bridges the gap between stochastic theory, numerical optimization, and autonomous system deployment.

3.1 Extended Uncertainty Modeling

The first objective of the STARGATE framework is to advance the modeling and integration of uncertainties within optimal control problems. Building on the Unscented Guidance technique, this approach seeks to extend its applicability to a broader class of uncertainty distributions and introduces an efficient sigma-point strategy aimed at mitigating the curse of dimensionality. The validity of the approach is verified on a representative Zermelo-type problem, introduced in [24] as part of a broader family of optimal control problems of particular interest for the current era of space exploration (from proximity operations around asteroids with uncertain gravitational properties, to trajectory correction maneuvers in interplanetary transfers). The formulation of the example problem is:

$$\mathbf{x} := (x, y) \in \mathbb{R}^2, \qquad \mathbf{u} \in \mathcal{U} := \{(u_1, u_2) \in \mathbb{R}^2 : u_1^2 + u_2^2 = 1\},$$
$$p \sim \mathcal{N}(\mu_p, \sigma_p^2)$$

$$(Z) \begin{cases} \text{Minimize} & J[\mathbf{x}(\cdot), \mathbf{u}(\cdot), t_f] := t_f \\ \text{Subject to} & \dot{x}(t) = p\, y^3(t) + u_1(t) \\ & \dot{y}(t) = u_2(t) \\ & (\mathbf{x}(t_0), t_0) = (2.25, 1, 0) \\ & x(t_f) = 0 \end{cases} \tag{1}$$

where $\mathbf{x}$ is the state vector, $\mathbf{u}$ is the control vector, p is the "wind" parameter in the Zermelo problem formulation, which follows a normal distribution with mean $\mu_p = 10$ and standard deviation $\sigma_p = 2$. Figure 3 shows results for problem 1 obtained with both the deterministic and stochastic optimal control approaches: consistently with the illustration in Figure 2, we observe that the deterministic solution in Figure 3a performs better under nominal conditions but fails catastrophically under the uncertain realization of the problem, while the stochastic solution in Figure 3b robustly satisfies the formulation of problem 1.

While the shown results are limited to a simplified application (i.e., single source of uncertainty following a normal distribution), future works will focus on fully illustrating the formulation of the uncertainty modeling approach, emphasizing its beneficial effects towards an efficient and accurate representation of stochastic optimal control problems. Ultimately, the capability to retain accuracy in the uncertainty representation while reducing the computational burden is fundamental to enable uncertainty-

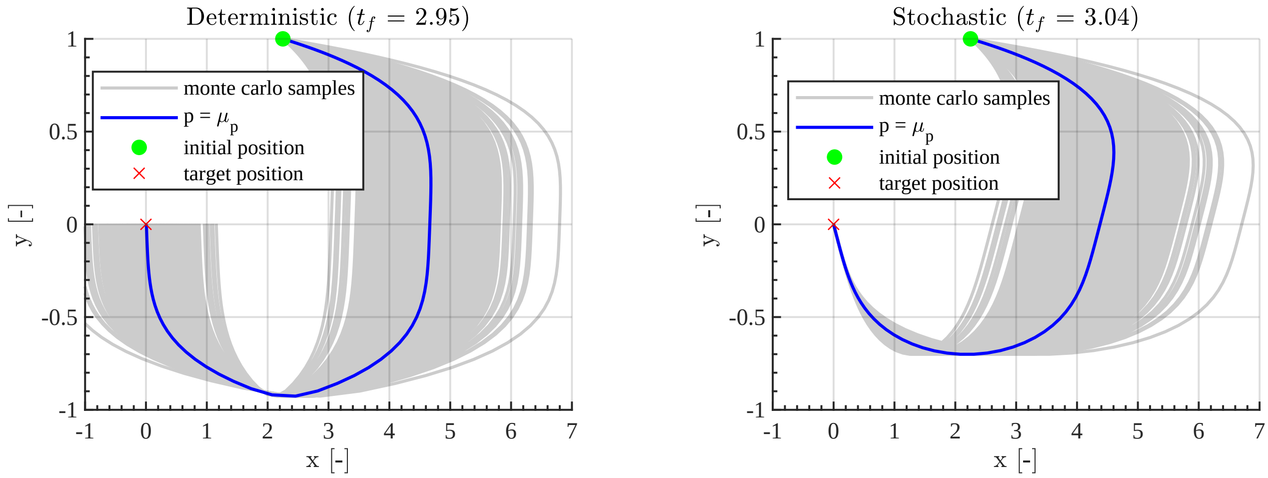

(a) Deterministic solution **(b) Stochastic solution**

Fig. 3 Comparison of the deterministic and stochastic optimal control approaches on problem Z, verified with 1000 Monte Carlo samples.

aware onboard-capable trajectory planning. Achieving this, however, requires solution strategies that can efficiently handle the high-dimensional and coupled nature of the resulting optimization problems.

3.2 Scalable Solution Technique

The second objective of the STARGATE framework focuses on the development of scalable and efficient solution techniques for stochastic optimal control problems. The problem resulting from the embedding of uncertainty (presented in subsection 3.1) suffers from both its high-dimensional nature and the inherent coupling among the differential equations evaluated at the different sampling points [26, 41]. The STARGATE framework aims to leverage the distributed nature of consensus-based optimization [49] to address such limitations.

The use of consensus optimization has raised concerns about its convergence accuracy and stability [51]. To address these issues, the STARGATE project proposes an iterative refinement strategy that blends principles from both direct and indirect optimal control. This hybrid approach enables sub-solutions to progressively converge toward optimality while preserving the decomposition benefits of consensus optimization. Combined with an innovative transcription, this procedure aims to deliver a computational framework capable of generating scalable, reliable, and onboard-feasible solutions.

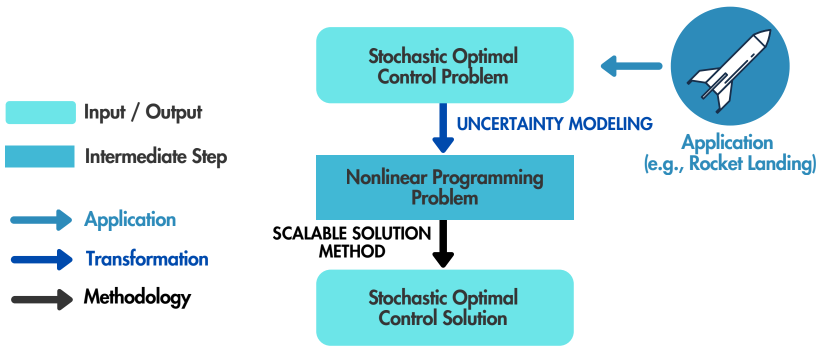

Fig. 4 STARGATE prototyped methodology.

The combination of efficient uncertainty modeling and scalable numerical solution methods yields the full prototyped STARGATE methodology, illustrated in Figure 4. The validation of such methodology is carried out on real autonomous systems, verifying the effectiveness of the approach when applied on real onboard processors.

3.3 Experimental Validation

The third objective of the STARGATE framework focuses on bridging theory and application by integrating the developed theoretical and numerical framework into a unified, deployable system. While simulation-based studies can effectively assess algorithmic performance, the implementation on representative flight processors enables a realistic assessment of its computational feasibility and robustness. This step, illustrated in Figure 5, is essential to demonstrate that uncertainty-aware trajectory optimization can be executed in real time within the constraints of modern space-qualified-like hardware.

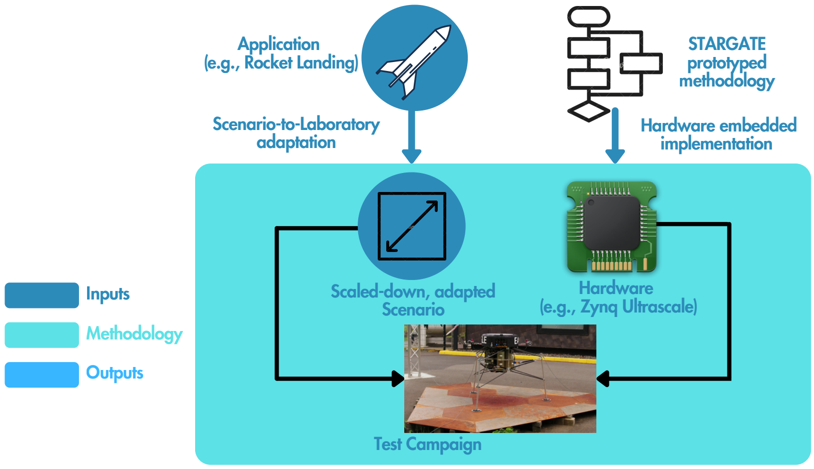

Fig. 5 STARGATE experimental validation setup.

Two complementary experimental test benches will be employed to validate the proposed framework across distinct operational regimes. The first focuses on a powered rocket landing scenario, representative of reusable launch vehicle descent operations. A suitable facility for this purpose will be selected to satisfy the requirements for the STARGATE framework (an example is the Environment for Autonomous GNC Landing Experiments (EAGLE) platform shown in Figure 6a for the reproduction of 6 Degrees of Freedom (DoF) rocket landing [52]). The defined setup shall provide a highly nonlinear and time-critical environment, ideal for evaluating the framework's capability to generate and update feasible trajectories in real time under significant dynamic uncertainties (e.g.: aerodynamic disturbances or propulsion mismatches.

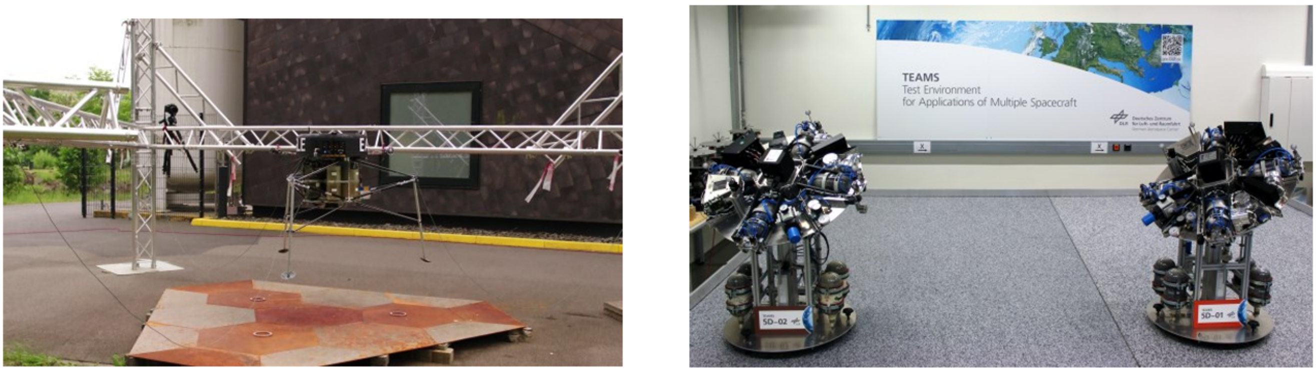

(a) EAGLE during a flight test.

(b) Test Environment for Applications of Multiple Spacecraft (TEAMS) laboratory.

Fig. 6 Potential facilities for the STARGATE experimental phase.

The second experimental scenario targets proximity operations, specifically rendezvous and docking maneuvers. A suitable facility for this class of missions will be identified to satisfy the desired test conditions (a potential candidate facility in that sense is the TEAMS shown in Figure 6b for 3- and 5-DoF rendezvous maneuvers). In this context, the STARGATE framework will be assessed in terms of guidance precision, robustness to modeling and injection errors, and capability to provide accurate trajectories despite uncertain modeling and sensor noise. Together, these two experimental campaigns constitute a comprehensive validation environment, spanning from fast, unstable flight dynamics to precise orbital proximity operations, demonstrating the versatility and onboard applicability of the proposed methodology.

3.4 Summary

The STARGATE framework introduces a unified approach to robust trajectory planning by embedding uncertainty directly within the optimization process rather than reacting to it through feedback control. This paradigm shift enables the generation of trajectories that are inherently resilient to stochastic disturbances, reducing the reliance on corrective control actions and improving both mission safety and efficiency. The ultimate goal of the STARGATE framework is to advance the state of the art of autonomous trajectory planning from the deterministic strategy in Figure 1 to the robust approach depicted in Figure 7.

Beyond its conceptual innovation, the STARGATE framework lays the foundation for scalable, real-time implementation of uncertainty-aware guidance algorithms onboard autonomous spacecraft. By combining advanced uncertainty modeling, consensus-based optimization, and hardware validation, the project aims to bridge the gap between theoretical robustness and practical autonomy. The resulting methodology is expected to enhance decision-making capabilities in safety-critical missions, paving the way toward a new generation of intelligent, self-reliant space systems.

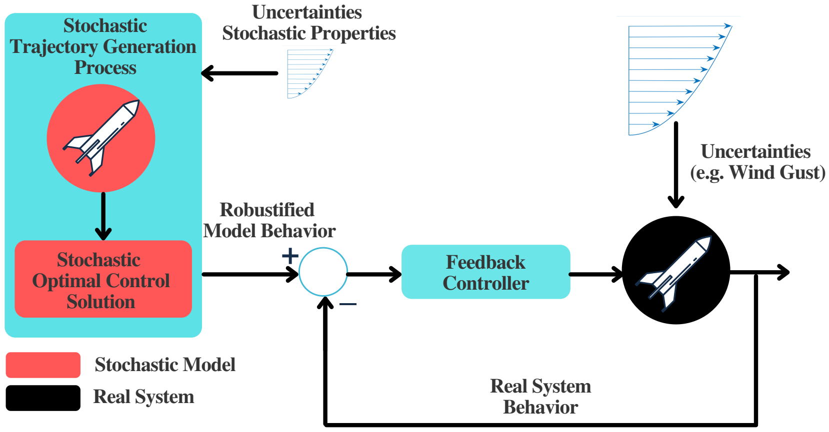

Fig. 7 STARGATE approach: uncertainty is embedded in the trajectory generation process.

4 Conclusions

This work presented the motivation, rationale, and conceptual design of the STARGATE framework, aiming at advancing the state of robust trajectory planning for autonomous space systems. The growing demand for autonomous decision-making to support the growth of the space sector towards increasingly complex and uncertain environments highlights the limitations of current approaches to trajectory planning. The use of deterministic models, where uncertainty is treated only as a disturbance rather than a key element of the design space, affects the safety and performance of current autonomous systems.

The proposed framework addresses this limitation by embedding uncertainty directly into the optimal control formulation, thus shifting from reactive to proactive robustness. Building upon and extending the principles of Unscented Guidance [24–31], the approach aims to introduce an efficient representation of different stochastic processes, which go beyond the classical Gaussian assumptions, into the formulation of optimal control problems. The adoption of consensus-based optimization alongside other elements of optimal control theory seeks to define a scalable and accurate solution process. Together, these developments enable the onboard generation of uncertainty-aware trajectories, and are envisioned for implementation and testing on embedded hardware in both rocket-landing and spacecraft rendezvous scenarios.

Future work, covering the period that goes from 2026 to 2030, will focus on the theoretical development and numerical implementation of the proposed techniques, defining the full methodology (illustrated in Figure 4) in terms of problem formulation and solution, ultimately followed by its integration with existing architectures and testing on physical systems. The continued development of the STARGATE framework is expected to provide key insights and to contribute to reduce the gap between optimal control theory and onboard autonomy, enabling safer, cheaper, and more adaptive space missions.

Acknowledgments

Funded by the European Union (ERC, STARGATE, 101169959). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

Funded by the European Union

European Research Council
Established by the European Commission

Declaration of Use of Artificial Intelligence

Artificial intelligence was not used in the work presented.

References

[1] P. Apkarian, D. Noll, and A. Falcoz. How we saved the rosetta mission. Technical report, Toulouse Institute of Mathematics, 2016.

[2] European Space Agency. Space environment statistics – objects in orbit. Space Debris User Portal.

[3] Ars Technica. The world just set a record for sending the most rockets into orbit, apr 2023. Retrieved on 01-Dec-2023. https://arstechnica.com/science/2022/01/thanks-to-china-and-spacex-the-world-set-an-orbital-launch-record-in-2021/#:~:text=The%20busy%20year%20in%202021,70%20orbital%20rockets%20a%20year.

[4] Ars Technica. Upper-level winds force a seven-week delay for the vega rocket, 2020. 7-Feb-2020. https://arstechnica.com/science/2020/07/launch-of-europes-smallest-rocket-is-delayed-seven-weeks-for-weather/.

[5] K. J. Astrom. *Introduction to Stochastic Control Theory*. Dover Publications, 2006. ISBN: 0486445313.

[6] W. Fleming. *Deterministic and Stochastic Optimal Control*. Springer, 1975. ISBN: 0387901558.

[7] D. Bertsekas. *Stochastic Optimal Control: the Discrete Time Case*. Athena Scientific, 2007. ISBN: 1886529035.

[8] M. Grant and M. A. Bolender. Rapid, robust trajectory design using indirect optimization methods. In *AIAA Atmospheric Flight Mechanics Conference*, 2015. doi: 10.2514/6.2015-2401.

[9] J. Ridderhof, J. Pilipovsky, and P. Tsiotras. Chance-constrained covariance control for low-thrust minimum-fuel trajectory optimization. In *AAS / AIAA Astrodynamic Specialist Conference*, 2020.

[10] K. Seywald and H. Seywald. Desensitized optimal control. In *AIAA Scitech Conference*, 2019. doi: 10.2514/6.2019-0651.

[11] Tommaso Robbiani, Marco Sagliano, Francesco Topputo, and Hans Seywald. Fast desensitized optimal control for rocket-powered descent and landing. *Journal of Guidance, Control, and Dynamics*, 48(11):2480–2494, 2025. doi: 10.2514/1.G009058.

[12] K. DeMars, R. H. Bishop, and M. K. Jah. Entropy-based approach for uncertainty propagation of nonlinear dynamical systems. *Journal of Guidance, Control, and Dynamics*, 36(4):1047–1057, 2014. doi: 10.2514/1.58987.

[13] V. Vittaldev, R. P. Russell, and R. Linares. Spacecraft uncertainty propagation using gaussian mixture models and polynomial chaos expansions. *Journal of Guidance, Control, and Dynamics*, 39(12):2615–2626, 2016. doi: 10.2514/1.G001571.

[14] V. Vittaldev and R. P. Russell. Multidirectional gaussian mixture models for nonlinear uncertainty propagation. *Computer Modeling in Engineering and Science*, 111:83–117, 2016. doi: 10.3970/cmes.2016.111.083.

[15] S. Boone and J. McMahon. Non-gaussian chance-constrained trajectory control using gaussian mixtures and risk allocation. In *IEEE 61st Conference on Decision and Control (CDC)*, 2022. doi: 10.1109/CDC51059.2022.9993274.

[16] S. Servadio, W. Parker, and R. Linares. Uncertainty propagation and filtering via the koopman operator in astrodynamics. *Journal of Spacecraft and Rockets*, 60(5):1639–1655, 2023. doi: 10.2514/1.A35688.

[17] M. Manzi and M. Vasile. Analysis of stochastic nearly-integrable dynamical systems using polynomial chaos expansions. In *AAS/AIAA Astrodynamics Specialist Conference*, 2020. doi: 10.1007/s10472-020-09723-1.

[18] M. Vasile. Practical uncertainty quantification in orbital mechanics. In *Satellite Dynamics and Space Missions*, pages 298–328. Springer, 2019. doi: 10.1007/978-3-030-20633-8_7.

[19] M. Vasile. Polynomial representation of model uncertainty in dynamical systems. In *Advances in Evolutionary and Deterministic Methods for Design, Optimization, and Control in Engineering and Sciences*, pages 419–432. Springer, 2019.

[20] J. Fisher and R. Bhattacharya. Optimal trajectory generation with probabilistic system uncertainty using polynomial chaos. *Journal of Dynamic Systems, Measurement, and Control*, 133(1), 2010. doi: 10.1115/1.4002705.

[21] L. J. Whittle and M. Sagliano. Stochastic optimal trajectory generation via multivariate polynomial chaos. In *2018 AIAA Guidance, Navigation, and Control Conference*. American Institute of Aeronautics and Astronautics, 2018. doi: 10.2514/6.2018-0849.

[22] D. Xiu. *Numerical Methods for Stochastic Computations: A Spectral Method Approach*. Princeton University Press, 2010. ISBN: 0691142122.

[23] X. Jiang, S. Li, R. Furfaro, Z. Wang, and J. Yuandong. High-dimensional uncertainty quantification for mars atmospheric entry using adaptive generalized polynomial chaos. *Aerospace Science and Technology*, 107, 2020. doi: 10.1016/j.ast.2020.106240.

[24] I. M. Ross, R. Proulx, and M. Karpenko. Unscented optimal control for orbital and proximity operations in an uncertain environment: A new zermelo problem. In *AIAA/AAS Astrodynamics Specialist Conference*. American Institute of Aeronautics and Astronautics, 2014. doi: 10.2514/6.2014-4423.

[25] D. Telen, M. Vallerio, L. Cabianca, B. Houska, J. V. Impe, and F. Logist. Approximate robust optimization of nonlinear systems under parametric uncertainty and process noise. *Journal of Process Control*, 33:140–154, 2015. doi: 10.1016/j.jprocont.2015.06.011.

[26] I. M. Ross, R. J. Proulx, M. Karpenko, and Q. Gong. Riemann–stieltjes optimal control problems for uncertain dynamic systems. *Journal of Guidance, Control, and Dynamics*, 38(7):1251–1263, 2015. doi: 10.2514/1.G000505.

[27] N. Ozaki, S. Campagnola, R. Funase, and C. H. Yam. Stochastic differential dynamic programming with unscented transform for low-thrust trajectory design. *Journal of Guidance, Control, and Dynamics*, 41(2):377–387, 2018. doi: 10.2514/1.G002367.

[28] N. Ozaki, S. Campagnola, and R. Funase. Tube stochastic optimal control for nonlinear constrained trajectory optimization problems. *Journal of Guidance, Control, and Dynamics*, 43(4):645–655, 2020. doi: 10.2514/1.G004363.

[29] C. Greco, S. Campagnola, and M. Vasile. Robust space trajectory design using belief optimal control. *Journal of Guidance, Control, and Dynamics*, 45(6):1060–1077, 2022. doi: 10.2514/1.G005704.

[30] D. González-Arribas, M. Soler, and M. Sanjurjo-Rivo. Robust aircraft trajectory planning under wind uncertainty using optimal control. *Journal of Guidance, Control, and Dynamics*, 41(3):673–688, 2018. doi: 10.2514/1.G002928.

[31] L. Ricciardi, C. Maddock, and M. Vasile. Robust trajectory optimisation of a tsto spaceplane using uncertainty-based atmospheric models. In *23rd AIAA International Space Planes and Hypersonic Systems and Technologies Conference*, 2020. doi: 10.2514/6.2020-2403.

[32] Jeffrey Uhlmann. *Dynamic Map Building and Localization: New Theoretical Foundations*. PhD thesis, University of Oxford, 1995.

[33] J. Crassidis. Sigma-point kalman filtering for integrated gps and inertial navigation. *IEEE Transactions on Aerospace and Electronic Systems*, 42(2):750–756, 2006. doi: 10.1109/TAES.2006.1642588.

[34] N. Adurthi, P. Singla, and T. Singh. The conjugate unscented transform – an approach to evaluate multidimensional expectation integrals. In *American Control Conference*, 2012. doi: 10.1109/ACC.2012.6314970.

[35] N. Adurthi, P. Singla, and T. Singh. Conjugate unscented transformation: Applications to estimation and control. *Journal of Dynamic Systems, Measurement, and Control*, 140(3):1–22, 2017. doi: 10.1115/1.4037783.

[36] D. Ebeigbe, T. Berry, M. Norton, A. J. Whalen, D. Simon, T. Sauer, and S. J. Schiff. A generalized unscented transform for probability distributions, 2021. Arxiv Preprint.

[37] Z. Stojanovski and D. Savransky. Higher-order unscented estimator. *Journal of Guidance, Control, and Dynamics*, 44(12):2186–2198, 2021. doi: 10.2514/1.G006109.

[38] T. Gerstner and M. Griebel. Numerical integration using sparse grids. *Numerical Algorithms*, 18:209–232, 2002. doi: 10.1023/A:1019129717644.

[39] P. E. Gill, W. Murray, and M. A. Saunders. User's guide for snopt version 7: Software for large-scale nonlinear programming. Technical report, University of California, San Diego, 2008.

[40] A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming*, 106:25–57, 2006. doi: 10.1007/s10107-004-0559-y.

[41] A. Selim and I. Ozkol. Multi-phase robust optimization of a hybrid guidance architecture for launch vehicles. In *10th EUCASS Conference*, 2023. doi: 10.13009/EUCASS2023-396.

[42] A. Selim and I. Ozkol. Robust trajectory optimization of constrained re-entry flight. In *10th International Conference on Recent Advances in Air and Space Technologies*, 2023. doi: 10.1109/RAST57548.2023.10197982.

[43] L. S. Pontryagin, V. G. Boltyanskii, R. V. Gamkrelidze, and E. F. Mishchenko. *The Mathematical Theory of Optimal Processes*. CRC Press, 1987. ISBN: 2881240771.

[44] Marco Sagliano, Ping Lu, David Seelbinder, and Stephan Theil. Analytical treatise on endo-atmospheric fuel-optimal rocket landings. *Journal of Guidance, Control, and Dynamics*, 48(3):450–469, 2025. doi: 10.2514/1.G008547.

[45] M. Sagliano. Performance analysis of linear and nonlinear techniques for automatic scaling of discretized control problems. *Operations Research Letters*, 42(3):213–216, 2014. doi: 10.1016/j.orl.2014.03.003.

[46] L. Huneker, M. Sagliano, and Y. Arslantas. Spartan: An improved global pseudospectral algorithm for high-fidelity entry-descent-landing guidance analysis. In *30th International Symposium on Space Technology and Science*, Kobe, Japan, 2015.

[47] M. Sagliano, S. Theil, V. D'Onofrio, and M. Bergsma. Spartan: A novel pseudospectral algorithm for entry, descent, and landing analysis. In *4th CEAS Eurognc Conference*, 2017. doi: 10.1007/978-3-319-65283-2_36.

[48] I. M. Ross. A universal birkhoff theory for fast trajectory optimization. ArXiv preprint, 2023.

[49] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2010. doi: 10.1561/2200000016.

[50] B. O'Donoghue, G. Stathopoulos, and S. Boyd. A splitting method for optimal control. *IEEE Transactions on Control Systems Technology*, 21(6):2432–2442, 2013. doi: 10.1109/TCST.2012.2231960.

[51] Y. Zheng, G. Fantuzzi, A. Papachristodoulou, P. Goulart, and A. Wynn. Chordal decomposition in operator splitting methods for sparse semidefinite programs. *Mathematical Programming*, 180:489–532, 2020. doi: 10.1007/s10107-019-01366-3.

[52] Michael Dumke, Guilherme Trigo, Marco Sagliano, Piyapat Saranrittichai, and Stephan Theil. Design, development, and flight testing of the vertical take-off and landing gnc testbed eagle. *CEAS Space Journal*, 08 2019. doi: 10.1007/s12567-019-00269-5.