(Preprint) AAS 25-610

INFORMATION-BASED GUIDANCE FOR ANGLES-ONLY RELATIVE NAVIGATION IN NONLINEAR DYNAMICS

Sergio Bonaccorsi[*], Alberto Fossà[†], Roberto Armellin[‡], Pierluigi Di Lizia[§], Renato Zanetti[¶]

The increasing presence of small spacecraft in environments characterized by nonlinear dynamics demands autonomous navigation strategies capable of dealing with limited sensing capabilities. This work presents a guidance policy that improves angles-only relative navigation performance through impulsive maneuvers that maximize information gain. The maneuver is optimized using both nonlinear programming and polynomial optimization under constraints on fuel consumption and tracking of a reference trajectory. The second method outperforms in terms of computational time, proving to be a valuable initial step for onboard guidance algorithms implementation. The test cases involving two spacecraft on close cislunar orbits demonstrate improved state observability. A sensitivity analysis provides insight into the behavior of the methods under uncertainties. Applications of such guidance policy may include formation flying, satellite inspection, and asteroid exploration missions, where enhanced navigation performance is essential for mission success in complex dynamical environments.

INTRODUCTION

A growing number of spacecraft is expected to operate in lunar and cislunar regimes. At the same time, space exploration is increasingly leveraging small platforms like SmallSats and CubeSats, which offer cost-effective opportunities for scientific and technological advancements. For both scientific and economic purposes, these small spacecraft generally carry an optical payload capable of acquiring images of target objects. Such data can then be processed through appropriate image processing pipelines to finally provide angular measurements of a given target. Issues related to state estimation operations in such a chaotic dynamics by relying on angles-only measurements are therefore becoming more and more relevant.[1] The primary weakness of this navigation method is that range information is not directly measured, thus leading to observability issues. Within this context, the motivations for enabling guidance and navigation co-design become evident. Employing control actions to enhance state estimation performance can be critically advantageous across a variety of mission scenarios, such as spacecraft formations, satellite inspection, and asteroid exploration. This approach allows the spacecraft to actively improve the information content of its

[*]PhD Candidate, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156 Milan, Italy.

[†]Postdoctoral Fellow, Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, Austin, TX, 78712.

[‡]Professor, Te Pūnaha Ātea – Space Institute, The University of Auckland, Auckland 1010, New Zealand

[§]Associate Professor, Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa 34, 20156 Milan, Italy.

[¶]Associate Professor, Department of Aerospace Engineering and Engineering Mechanics, The University of Texas at Austin, Austin, TX, 78712.

measurements, thereby improving navigation accuracy in situations where observability is otherwise limited. Moreover, as autonomous spacecraft operations increasingly assume a central role in space mission design, the interest for computationally efficient algorithms to be used for real-time onboard operations is growing.

Information-based guidance has emerged as a promising framework for enhancing navigation performance in angles-only scenarios where observability is limited. A key contribution is the work of Woffinden and Geller, who formally analyze the observability conditions of angles-only navigation systems.[2] In a two-body dynamics framework, they show that without control input, the system is unobservable due to the range ambiguity in the relative state. However, they demonstrate that a properly designed impulsive maneuver can ensure full system observability, highlighting the crucial role of control actions in enabling reliable state estimation from bearing measurements alone. More recently, Greaves and Scheeres have extended these principles to inter-satellite tracking in cislunar space. They present an analytical formulation for the computation of impulsive maneuvers that maximize the deviation in optical-only measurements, known as the Maximum Measurement Deviation (MMD) maneuver. The approach relies on linearized relative dynamics and introduces an eigenvalue-based solution to identify the control direction that maximizes the sensitivity of the measurement space to changes in the state. The MMD maneuver is elegant and computationally efficient, but it does not explicitly account for practical mission constraints such as reference trajectory adherence. As such, it can be used as a foundational benchmark for more general guidance strategies.

This study aims to evaluate the potential and limitations of an information-driven guidance policy. The idea is to address a guidance problem focused on maximizing the information content of angles-only measurements by pursuing the relative range uncertainty reduction as proposed by Greaves and Scheeres.[3] Starting from this foundation, the study introduces incremental improvements over prior approaches by incorporating key operational constraints. These include maintaining proximity to a reference trajectory and keeping fuel consumption below a specified threshold. The proposed solution consists of an impulsive maneuver designed to reduce the relative uncertainty, thereby improving overall mission performance and reducing the need for future intensive navigation efforts. Furthermore, a novel contribution is offered through the formulation of the guidance problem within a polynomial optimization framework.

The statement of the problem is first introduced, with details about the formulation in two different frameworks, namely the nonlinear programming (NLP) problem and the polynomial optimization problem (POP). The numerical results obtained from these two approaches applied on two different scenarios are then presented and compared in terms of performance through different metrics. A sensitivity analysis is performed to assess the robustness of the solutions under uncertainties and in different configurations. Key observations on the results are then discussed. Finally, the conclusions summarize the main findings, highlight the objectives achieved, and propose directions for future research.

PROBLEM STATEMENT

In the context of the Circular Restricted Three Body Problem (CR3BP),[4] the following guidance problem is considered for the design of information-driven maneuvers. The objective is to minimize the relative range uncertainty between the observer spacecraft and a target in a fixed time interval. Additionally, nonlinear constraints on a limited fuel budget and maximum distance from the reference trajectory are included. The timeline of the analyzed scenario is illustrated in Figure 1. The red bar indicates the relative navigation phase, while the green arrow marks the moment at which the impulsive maneuver, derived from the optimization process, is applied. The days elapsed since the starting point are shown above the timeline, whereas the relevant time markers used to define the quantities involved in the optimization problem are displayed below.

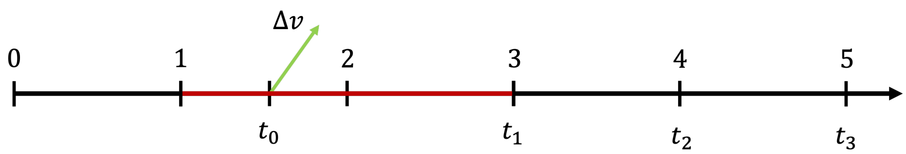

Figure 1: Timeline of the guidance problem. The red bar represents the relative navigation phase and the green arrow marks the maneuvering point. Above the timeline, the days after the starting point are reported. Below the timeline, the times used within the optimization problem are reported.

Nonlinear Programming Problem

Following the work by Greaves and Scheeres,[3] the same quantity representing the relative range uncertainty will be considered for the optimization problem.

$$\sigma_\rho^2(\Delta\boldsymbol{v}) = \mathbf{M}_\rho^T \tilde{\mathbf{P}}_f \mathbf{M}_\rho \tag{1}$$

Here, $\tilde{\mathbf{P}}_f$ is the covariance at the final time of the navigation window, and $\mathbf{M}_\rho$ maps the covariance onto the relative range direction between the observer and the target $\hat{\boldsymbol{\rho}}_f$, such that

$$M_\rho = \begin{bmatrix} \mathbf{0}_{6\times1}^T & \hat{\boldsymbol{\rho}}_f^T & \mathbf{0}_{3\times1}^T \end{bmatrix}^T \tag{2}$$

In the NLP framework, the problem can be then formulated as follows:

$$\min_{\Delta\boldsymbol{v}} \sigma_\rho^2(\Delta\boldsymbol{v}) \tag{3}$$

subject to the equality constraints:

$$\boldsymbol{x}(t_1) = \boldsymbol{x}(t_0) + \int_{t_0}^{t_1} \boldsymbol{f}(\boldsymbol{x}(t), t)\,\mathrm{d}t \tag{4a}$$

$$\boldsymbol{x}(t_0) = \boldsymbol{x}_0^- + \begin{bmatrix} \mathbf{0}_{n\times1} \\ \Delta\boldsymbol{v} \end{bmatrix} \tag{4b}$$

and inequality constraints:

$$[\boldsymbol{x}(t_k) - \bar{\boldsymbol{x}}_k]^T \boldsymbol{W} [\boldsymbol{x}(t_k) - \bar{\boldsymbol{x}}_k] - 1 \leq 0 \tag{5a}$$

$$\Delta\boldsymbol{v}^T \Delta\boldsymbol{v} - \Delta v_{\text{max}}^2 \leq 0 \tag{5b}$$

where σ_ρ^2 is the uncertainty associated to the relative range ρ, $\boldsymbol{f}(\boldsymbol{x}(t), t)$ is the ballistic dynamics, $\boldsymbol{x}_0^- \in \mathbb{R}^6$ is the initial state before the maneuver, $\bar{\boldsymbol{x}}_k \in \mathbb{R}^6$ is the reference state at time t_k, $\Delta\boldsymbol{v} \in \mathbb{R}^3$ is the impulsive maneuver, $\boldsymbol{W} \in \mathbb{R}^{3\times3}$ is a weight matrix chosen to limit the distance from the reference trajectory. Then, t_k are the k-th times represented in Figure 1, where t_0 is the maneuvering time, t_1 is where both the cost function and the first nonlinear constraint are evaluated, and t_2 and t_3 are the second and third evaluation points of the constraint. The presence of more than a single time for the evaluation of the constraint on the distance from the reference path is chosen such that a sort of continuous bounding along the trajectory is imposed. The time for the evaluation of the cost function t_1 must coincide with the end of the navigation window.

Equations (3) to (5) define an NLP problem that can be readily solved with dedicated algorithms such as MATLAB `fmincon`.

Polynomial Optimization Problem

According to the needs of reducing the computational time for a possible onboard implementation, the numerical integration of the dynamics at each iteration should be avoided. An efficient alternative is to replace the integration with a high-order Taylor expansion of $\boldsymbol{x}(t_1)$ as a function of $\Delta\boldsymbol{v}$, computed via differential algebra (DA).

Polynomial expansion of the constraints. Consider the spacecraft state $\boldsymbol{x}$ defined as:

$$\boldsymbol{x} = \left[\boldsymbol{r}^T \;\; \boldsymbol{v}^T\right]^T = [x \;\; y \;\; z \;\; v_x \;\; v_y \;\; v_z]^T \in \mathbb{R}^{6\times 1} \tag{6}$$

and the impulsive maneuver $\Delta\boldsymbol{v}$ given by:

$$\Delta\boldsymbol{v} = [\Delta v_x \;\; \Delta v_y \;\; \Delta v_z]^T \in \mathbb{R}^{3\times 1} \tag{7}$$

A polynomial expansion of the final state can be efficiently obtained using DA techniques.[5] The impulsive maneuver is firstly initialized as:

$$[\Delta\boldsymbol{v}] = \boldsymbol{0}_{3\times 1} + \delta\boldsymbol{v} \tag{8}$$

where the square brackets denote Taylor polynomials, and $\delta\boldsymbol{v} = [\delta v_x \;\; \delta v_y \;\; \delta v_z]^T$ are the three independent DA variables. Then, evaluating Eq. (2) in the DA framework results in the following expression for the final state:

$$[\boldsymbol{x}(t_1)] = \mathcal{T}^{(n)}_{\boldsymbol{x}(t_1)}(\delta\boldsymbol{v}) \tag{9}$$

which is a vector of n-th order Taylor polynomials in $\delta\boldsymbol{v}$. Substituting Equation (9) for different propagation times t_k into the right-hand side of Equation (5a) finally yields the following approximation of the nonlinear constraint:

$$\begin{aligned}\left[d^2_{\text{ref}}(t_k)\right] &= \{[\boldsymbol{x}(t_k)] - \bar{\boldsymbol{x}}_k\}^T \boldsymbol{W} \{[\boldsymbol{x}(t_k)] - \bar{\boldsymbol{x}}_k\} \\ &= \mathcal{T}^{(n)}_{d^2_{\text{ref}}(t_k)}(\delta\boldsymbol{v})\end{aligned} \tag{10}$$

which is again a n-th order Taylor polynomial in $\delta\boldsymbol{v}$.

Optimization as a nonlinear programming problem. Given Equation (10), the original optimization problem is recast into the following POP:

$$\min_{\Delta\boldsymbol{v}} \sigma^2_\rho(\Delta\boldsymbol{v}) \tag{11}$$

subject to:

$$T^{(n)}_{d^2_{\text{ref}}(t_k)}(\Delta v) - 1 \leq 0 \tag{12a}$$

$$\Delta\boldsymbol{v}^T\Delta\boldsymbol{v} - \Delta v^2_{\text{max}} \leq 0 \tag{12b}$$

The constraint in Equation (12a) substitutes both Equations (4) and (5a) using a k-th order approximation of the dynamics and of the squared distance from the reference trajectory. As POPs are a particular case of NLPs, this problem can readily be solved using a generic NLP solver.

Sum-of-squares optimization. Although POPs can be solved using generic NLP solvers, these algorithms treat the objective and constraints functions as black-boxes, and do not exploit their polynomial structure. Moreover, most solvers are gradient-based, meaning that the provided solution might be only locally optimal. In contrast, Sum-of-squares (SOS) optimization[6,7] exploits the polynomial structure of the problem to find a lower bound of the cost function over the feasible set. The global optimum (or optima) are then retrieved as the points where this bound is attained.

To solve the POP defined by Equations (11) and (12) using SOS optimization, an additional scalar variable $\alpha \in \mathbb{R}$ is introduced. This variable represents the sought-after lower bound of the cost function. The optimization problem is then reformulated as:

$$\max_{\alpha, \Delta v} \alpha \tag{13}$$

subject to the inequality constraints:

$$\alpha - \sigma_\rho^2(\Delta \boldsymbol{v}) \leq 0 \tag{14a}$$

$$T^{(n)}_{d^2_{\text{ref}}(t_k)}(\Delta v) - 1 \leq 0 \tag{14b}$$

$$\Delta \boldsymbol{v}^T \Delta \boldsymbol{v} - \Delta v^2_{\text{max}} \leq 0 \tag{14c}$$

SOS optimization builds a hierarchy of relaxed problems that can be solved efficiently via semidefinite programming (SDP).[7] Their solutions provide increasingly tighter lower bounds of the cost function in Equation (13) until the optimal value α^* is attained. The optimal solution to the original POP is then recovered from the underlying relaxation.

NUMERICAL RESULTS

The goal of the following results is to show that having a specific information-based guidance policy allows to reach greater advantages in the uncertainty reduction, while keeping the fuel consumption and the reference tracking constraints satisfied. For this reason, the performance of the maneuvers coming from the resolution of the NLP problem and the POP are compared against the Target Point Approach (TPA), which is a typical method for the design of station-keeping maneuvers. The comparison is done because the goal is to reduce the relative uncertainty by exploiting the change of geometry generated by the impulsive maneuver, but such a change, and the subsequent improvement, can be achieved with any applied maneuver. Therefore, the idea is to highlight the benefits of the previously introduced guidance policy for this kind applications.

These analyses are indeed of interest because the most frequently executed maneuver throughout the mission is typically the station-keeping one, which ensures the spacecraft remains close to its reference trajectory. Consequently, if a reduction of the relative range uncertainty with respect to a target is required, mission operators may consider scheduling a navigation window that leverages the change in relative geometry induced by the already-planned station-keeping maneuver to enhance estimation performance. In such cases, introducing an additional, dedicated information-driven maneuver might not be desired unless the expected improvement in estimation performance clearly justifies the added complexity and resource expenditure.

Two scenarios are considered to evaluate the algorithms under different dynamical conditions. In both cases, the impulsive maneuver is scheduled at apolune to minimize the risk of significant deviation from the reference trajectory.

The time evolution of the relative range uncertainty is computed according to a linear covariance analysis around the true states. This is equivalent to a Cramer-Rao lower bound analysis, which, for nonlinear systems with Gaussian process and measurement noise, is generated through the Extended Kalman Filter (EKF) covariance prediction and correction equations, linearized around the true system state.

Short-Term Scenario

In this first test case, the observer spacecraft follows a Near-Rectilinear Halo Orbit (NRHO) trajectory, while the target follows a very close quasi-periodic orbit designed to share the same orbital period, thereby inducing a naturally bounded relative motion. The results here presented relies on the short timeline of Figure 1. Indeed, in Figure 3 the blue band represents the relative navigation window in which the covariance is updated through the linearized approach. The values of the thresholds of the nonlinear constraints are of 100 km for the radius of the reference tracking and 1 m/s for the fuel budget.

Before presenting the optimization results, Figure 2 illustrates that any applied maneuver enhances estimation performance. To generate this plot, 100 randomly oriented impulsive maneuver vectors, each with a 1 m/s magnitude, were propagated through the same pipeline. The corresponding uncertainty evolutions appear as gray curves. The optimal solutions, obtained by solving the NLP and the POP, are highlighted in red and yellow, respectively. This comparison underlines both the beneficial effect of applying a maneuver and the added value of specifically optimized maneuvers over arbitrary ones.

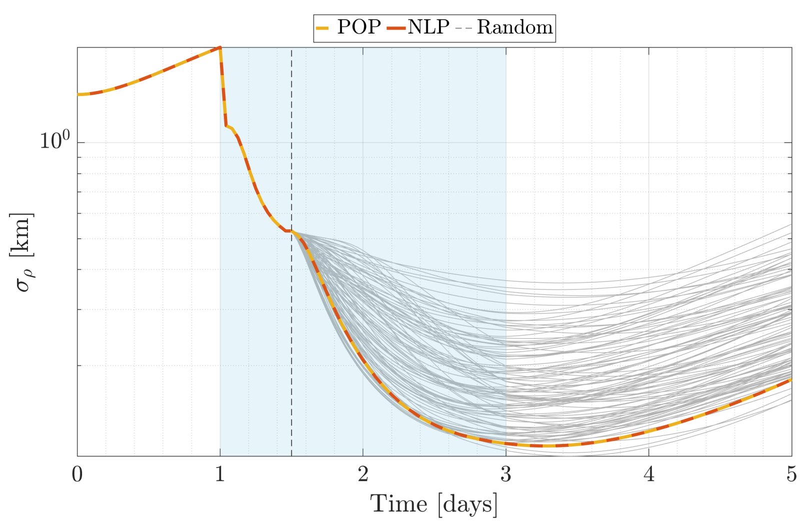

Figure 2: Comparison between state estimation performance generated by randomly oriented impulsive maneuver vectors and optimized solutions by the NLP and POP.

Concerning the optimization results against the TPA, from Figure 3 two main aspects are immediately evident. The first is that the optimal solutions given by NLP and POP reach a higher reduction of the relative range uncertainty. Second, these two solutions are superimposed, meaning that the resolution process within the polynomial framework agree with the fully nonlinear procedure.

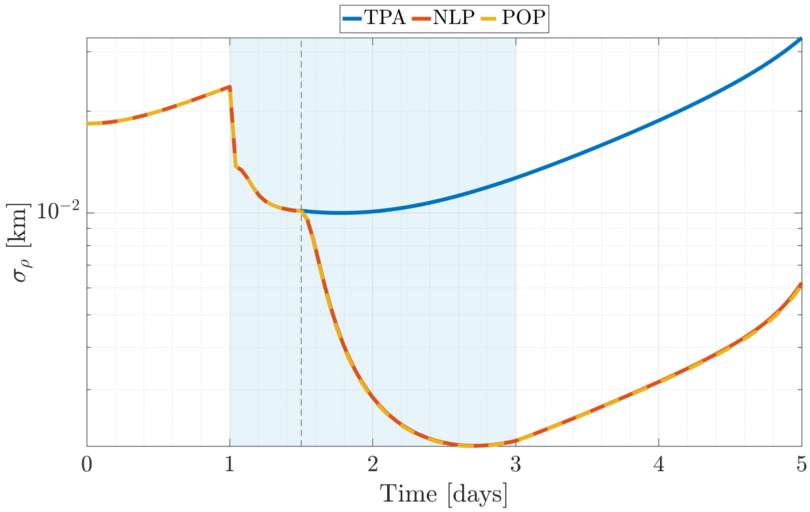

Figure 3: Comparison of the time evolution of the relative range uncertainty resulting from the application of the maneuvers coming from the TPA, the NLP, and the POP.

Other relevant peculiarities of the solutions are the fuel consumption and the ability to track the reference. Regarding the first, as many works from the literature have shown, the opposing nature between observability performance and fuel expenses results in a Pareto front, reflecting the trade-off between competing criteria. The existence of such competitive nature translate into the fact that imposing a larger fuel budget in the setup of the algorithm results in a higher reduction of the uncertainty level. This is clearly visible in Figure 3 considering the fuel consumption shown in Table 1.

The possibility of arbitrarily setting the fuel budget is a great advantage from an operational point of view, since the user can balance a trade-off between the fuel capacity of the spacecraft and the current need for an improvement in the navigation performance. However, it must be reminded that a lower bound in the level of estimated uncertainty exists, meaning that there is a limit in the achievable improvement. In addition, by running multiple simulations, it has been noted that the algorithm is always capable of reaching a lower value of relative uncertainty if the fuel expense is larger or equal to the one of the TPA solution. This means that the primary goal of the guidance policy has been achieved and that the user can decide to tune the fuel budget depending on the current needs. On the other hand, a larger fuel budget might translate into a larger separation distance from the reference trajectory, but always between the bounds of the imposed constraint.

Method	Δv
TPA	0.017 m/s
NLP	1.000 m/s
POP	1.000 m/s

Table 1: Fuel consumption for the TPA, NLP and POP solutions.

Long-Term Scenario

After verifying the effectiveness of the algorithms in producing solutions that reduce estimation uncertainty, the focus shifts toward assessing the long-term impact of applying an information-driven guidance policy. While the primary objective remains the immediate reduction of relative range uncertainty with a prescribed deviation from the nominal trajectory, it is equally important to understand how this improvement evolves over time and to quantify the overall benefit of the strategy. To this end, three key aspects are considered: the long-term evolution of the relative range uncertainty, the cumulative fuel consumption, and the reference tracking performance.

In this case, the observer spacecraft is on the halo orbit of the LUMIO mission, and the target follows a quasi-periodic orbit with matching period. The timeline depicted in Figure 4 illustrates the structure of this analysis. In the figure, blue bars indicate the relative navigation phases during which the covariance is updated for the three methods, while the orange bar corresponds to the phase where the covariance is updated only for the TPA solution. The green arrow marks the maneuvering point where the optimal solutions are applied (only for NLP and POP cases), and the red arrows indicate the subsequent station-keeping maneuvers. This means that, also for the optimal strategies, only the first maneuver is the result of an optimization process, while the other two are computed according to the TPA method. The number of days elapsed since the initial reference epoch is also shown above the timeline to provide temporal context.

From a practical standpoint, the objective is to assess whether the information gained through the application of optimal maneuvers is sufficient to skip the subsequent navigation window while still maintaining acceptable estimation performance. To evaluate this, the results are compared against the TPA strategy, for which all navigation windows are consistently exploited.

Figure 4: Timeline of the long-term scenario. The blue bars represent the relative navigation phases in which the covariance is updated for the three methods, while in the orange bar the covariance is updated only for the TPA. The green arrow marks the maneuvering point in which optimal solutions are applied (only for NLP and POP), the red arrows mark the next station-keeping maneuvers. Above the timeline, the days after the starting point are reported.

Method	Δv
TPA	0.9784 m/s
NLP	1.4738 m/s
POP	1.4894 m/s

Table 2: Fuel consumption for the TPA, NLP and POP solutions for the analysis on the long-term scenario.

The results shown in Figure 5 clearly indicate that, even when one navigation window is skipped, the uncertainty level resulting from the application of the optimal maneuvers remains comparable to that of the TPA strategy, which instead relies on all available navigation windows. This simple analysis gives a first hint on the potential effectiveness and robustness of the information-driven guidance approach in sustaining estimation performance over extended time horizons.

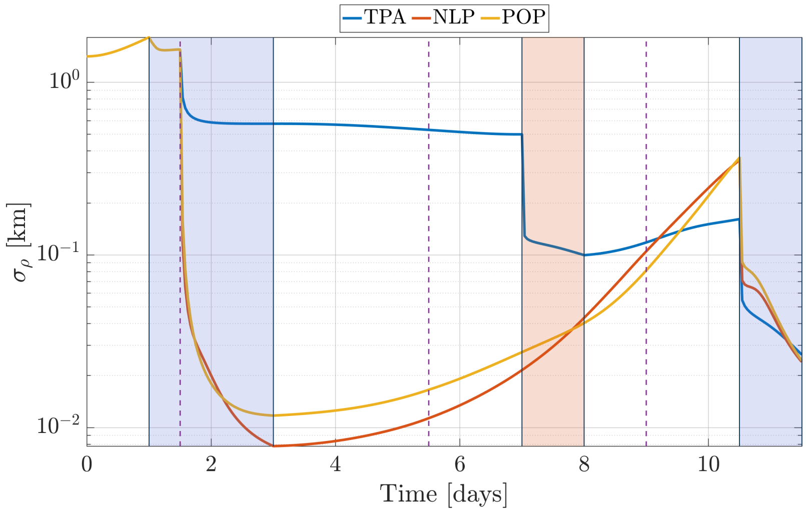

Figure 5: Comparison of the time evolution of the relative range uncertainty resulting from the application of the maneuvers coming from the TPA, the NLP, and the POP.

These considerations are also supported by the fuel consumptions reported in Table 2, which shows that, as expected, the TPA strategy results in the lowest average cost. This is consistent with its design objective of minimum fuel. However, this lower consumption comes at the cost of requiring frequent navigation updates to sustain acceptable estimation performance. In contrast, the NLP and POP strategies show slightly higher fuel consumption, which is the direct consequence of maneuvers designed to enhance the observability of the system. The result is a moderate increase in fuel usage in exchange for greater autonomy.

Accuracy of the Polynomial Approximation. From the results shown in Figure 5 and Table 2, it is evident that the fully nonlinear approach of the NLP and the polynomial-based framework of the POP produce different results, which is in contrast with the results presented in Figure 3, where the two optimized solutions coincide.

When solving the problem within a polynomial framework there are two main aspects to be monitored in order to assess the goodness of the solutions. These are the polynomial approximation of both the cost function and the nonlinear constraints, and the use of a specific solver, which can lead to sub-optimal solutions, not aligned with the NLP ones.

A deep analysis has been conducted to highlight the differences in the solutions obtained with three different resolution approaches, namely the pure NLP based on the nonlinear representation of both the cost function and the constraints, then the NLP based on the use of a polynomial approximation of the latter, which will be denoted as PolNLP, and finally, the use of a dedicated SOS optimization framework to solve the POP.

Different setup for the nonlinear constraints have been tested to assess their role in the resolution process. The results have shown that, for the NRHO case, all the approaches converge to the optimal solution for each setup of the thresholds for the constraints. For the halo orbit case instead, the it has been found that the optimal solution is achieved only with the NLP, while the PolNLP and the POP converge to the same sub-optimal solution.

These results indicate that the solver itself does not play a decisive role in the resolution process, as the PolNLP and POP methods yield consistent solutions. Instead, the critical factor is the polynomial approximation of the cost function and nonlinear constraints, which facilitates the problem's tractability and ensures convergence. However, this approximation may also introduce sub-optimality or limit the accuracy of the solution, particularly in scenarios where the nonlinearities are significant.

Degree of Nonlinearity Assessment. The cause of the difference among the reported solutions may therefore be assigned to the degree of nonlinearity of the problem at hand. It might be possible that the halo orbit case presents a higher degree of nonlinearity, which makes the polynomial approximation weaker and makes therefore converge the algorithms to a sub-optimal solution.

To better investigate this point, a quantitative analysis on the degree of nonlinearity between the two cases is performed. For this purpose, two indexes from Junkins and Singla[8] are used and here recalled.

$$\nu_{\Phi}(t) = \frac{\|\boldsymbol{x}(t) - \boldsymbol{\Phi}(t, t_0)\boldsymbol{x}_0\|}{\|\boldsymbol{x}(t)\|} \tag{15}$$

$$\nu_{\dot{x}}(t) = \frac{\|\boldsymbol{f}(\boldsymbol{x}, t) - \boldsymbol{A}\boldsymbol{x}\|}{\|\boldsymbol{f}(\boldsymbol{x}, t)\|} \tag{16}$$

where $\boldsymbol{x}(t)$ is the propagated reference state at time t, $\boldsymbol{\Phi}(t, t_0)$ is the state transition matrix from t_0 to t, $\boldsymbol{f}(\boldsymbol{x}, t)$ is the right-hand side of the CR3BP dynamics, and $\boldsymbol{A}$ is its Jacobian.

These indexes are evaluated along time within the interval of the simulation and the percentage of the epochs for which the halo dynamics has a higher degree of nonlinearity over the NRHO one is reported for both indexes in Table 3. This shows that the halo case presents stronger nonlinearities, which might reasonably be the cause for the lack of accuracy of the polynomial approximation.

Index	Percentage
$\nu_{\Phi}(t)$	81.7%
$\nu_{\dot{x}}(t)$	89.3%

Table 3: Percentage of the epochs for which the degree of nonlinearity of the halo is larger than the NRHO.

SENSITIVITY ANALYSIS

For a deep investigation on the behavior of the algorithms and the resulting performance of the associated solutions, a sensitivity analysis is performed. A large initial dispersion, with a 60 km value for the standard deviation in position and 10 cm/s for the velocity is selected to depict a situation where the satellite operates in the vicinity of the reference path but within a broader tolerance. The underlying rationale is that, by allowing the spacecraft to move freely within a larger yet constrained region around the nominal path, rather than enforcing strict adherence to a specific trajectory, the guidance problem gains additional degrees of freedom. This flexibility enables the identification of more favorable trajectories from an observability standpoint, thus leading to improved relative navigation performance through enhanced measurement geometry.

Method	Δv
TPA mean	0.2988 m/s
NLP mean	0.9456 m/s
POP mean	0.8469 m/s

Table 4: Average fuel consumption for the TPA, NLP, and POP solutions based on 600 MC samples.

For the halo scenario in the short-term timeline, two different cases will be considered. The first is intended to test the average performance of the methods in nominal conditions starting from 600 Monte Carlo (MC) samples extracted from the initial dispersion. The second analysis tests the effectiveness of the algorithms under uncertainties.

For both cases, the values of the thresholds of the nonlinear constraints are of 200 km for the reference tracking and 2.5 m/s for the fuel budget. These values are chosen to effectively highlight both the potential and the limitations of the information-driven guidance policy. In particular, the reference tracking bound is sufficiently large, and the fuel budget is high enough to permit maneuvers that can meaningfully enhance observability. This configuration enables a fair evaluation of how the guidance strategy exploits the available freedom to improve navigation performance.

Nominal Setup

Several observations can be made based on the results shown in Figure 6. First, a significantly large dispersion is observed in the time evolution of the relative range uncertainty across the Monte Carlo realizations. This variability is a direct consequence of the large initial state dispersion. As a result, the impact of the applied maneuvers on the estimation performance might relevantly vary across different samples.

There is a difference between the NLP and POP solutions, which indicates that the polynomial approximation used in the POP affects the resolution process, possibly driving the optimization toward sub-optimal maneuver directions, particularly in highly nonlinear regions of the solution space.

Both NLP and POP yield superior performance in terms of uncertainty reduction when compared to the TPA strategy. This confirms the advantage of an information-driven guidance approach in identifying maneuvers that effectively enhance observability.

As expected, the TPA strategy consistently delivers the best results in terms of reference tracking and fuel consumption, as seen in Table 4.

In summary, the analysis highlights the complementary strengths of the evaluated strategies. The information-driven approaches, namely NLP and POP, prove effective in enhancing state estimation performance, particularly in scenarios where flexibility in trajectory deviation is accepted. However, the POP method may still be prone to sub-optimal solutions due to the limitations of polynomial approximation in highly nonlinear regimes. On the other hand, the TPA strategy remains the best option when strict reference tracking is preferred. Therefore, the choice of the strategy should be guided by the mission priorities. If improving navigation accuracy is critical and some deviation from the reference path is acceptable, an optimal information-driven maneuver policy might be adopted. On the other hand, if adherence to the nominal orbit is the main driver, the TPA approach, or other well-known station-keeping strategies, offer more suitable solutions.

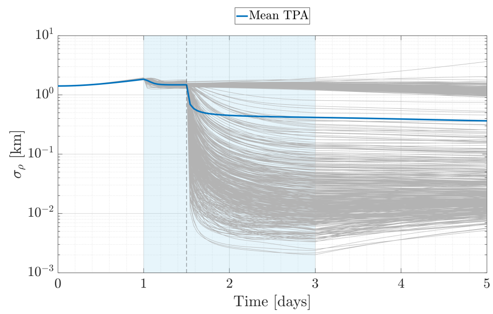

(a) TPA solution.

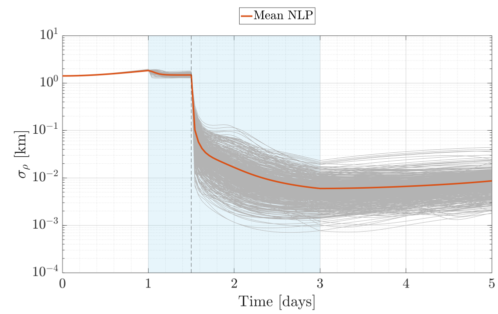

(b) NLP solution.

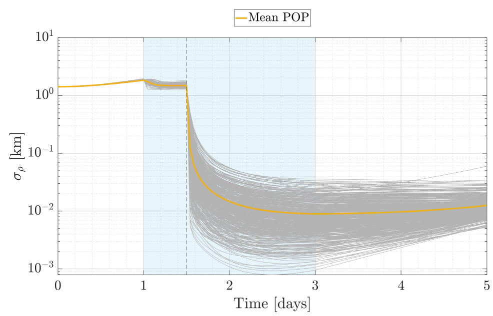

(c) POP solution.

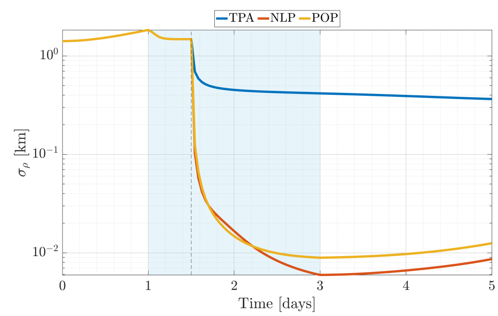

(d) Comparison between solutions.

Figure 6: halo case: time evolution of the relative range uncertainty based on 600 MC samples.

Performance Under Uncertainties

Here the behavior of the algorithm is evaluated when multiple sources of discrepancy from the nominal conditions are considered. In particular, the following aspects are included in the simulations: addition of an orbit determination error on the spacecraft state given as input to the algorithm, addition of a maneuver execution error, and the computation of the time evolution of the covariance through a navigation filter.

First, the algorithm is fed with an estimate of the true spacecraft state, and the resulting impulsive maneuver is instead applied to the true state. This procedure is consistent with a realistic mission scenario, where the computation of correction maneuvers is performed relying on the orbital knowledge provided by a previous orbit determination process. Consequently, the resulting computed maneuver is applied to the actual spacecraft state, introducing a discrepancy due to estimation errors and resulting in an imperfect realization of the intended correction strategy.

Then, when maneuver execution errors are considered, the resulting impulse has slightly different magnitude and direction. This is also a crucial factor to be considered in order to assess the goodness of the optimal solutions within a realistic framework.

Finally, a relevant layer of realism is added to the results previously shown. Up to this point, the Cramer-Rao lower bound was evaluated. Therefore, during the simulated navigation window, the spacecraft state is not estimated, while the uncertainty is reduced according to the value of the Kalman gain and the information contained in the Jacobian of the angles-only measurement model. Now, both the state and the covariance are predicted and corrected within an EKF. This implies that the results are affected by the presence of noise in the simulated measurements, and by the choice of some filter parameters such as the initial guess and the noise matrix.

The performance analyzed through this new setting should be considered more reliable, given that, in a real scenario, the time evolution of the covariance is actually produced through the operations of a navigation filter. However, it is important to note that the results and accuracy of such a filter depend on numerous factors and parameter choices. For this reason, only preliminary and general considerations will be presented in the following analysis, without attempting to capture the full complexity of an operational navigation system.

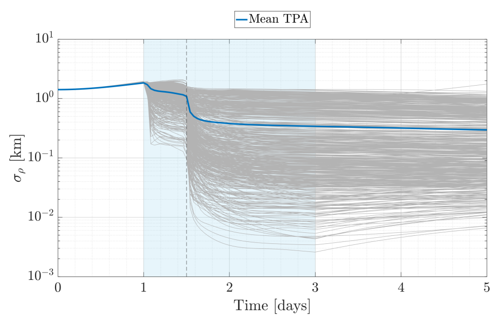

(a) TPA solution.

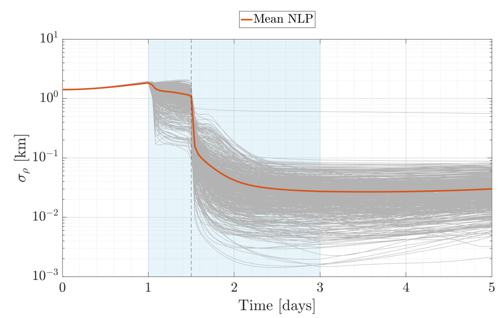

(b) NLP solution.

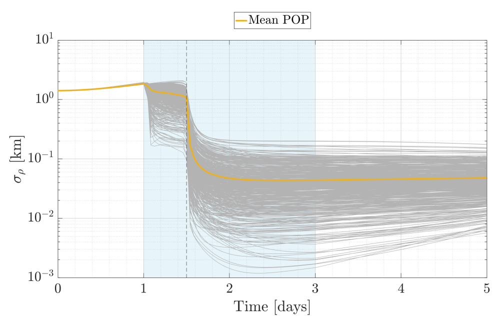

(c) POP solution.

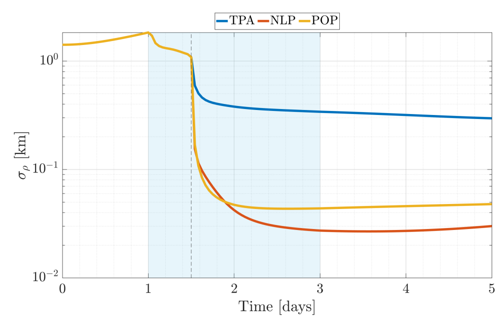

(d) Comparison between solutions.

Figure 7: Time evolution of the relative range uncertainty for the scenario with uncertainties based on 600 MC samples.

Method	Δv
TPA mean	0.3000 m/s
NLP mean	0.9435 m/s
POP mean	0.8490 m/s

Table 5: Average fuel consumption for the TPA, NLP and POP solutions for the scenario with uncertainties based on 600 MC samples.

First, the most relevant thing to notice in Figure 7 is that the role of the maneuver in improving the estimation performance is still present, even under the operations of a navigation filter. This confirms again the validity of an observability-driven guidance policy.

Then, for some MC samples we can notice a drop in the level of uncertainty at the beginning of the simulated navigation window. This may be justified by an overconfident behavior of the filter, which is initialized with a very good initial guess and provided with measurements affected by a very low level of noise. As said earlier, a deep analysis of these aspects is not the main goal. However, this confirms that the actual performance must be assessed through a complete state estimation algorithm to better highlight possible realistic outcomes.

According to this, another important aspect is the presence of certain MC realizations in which the maneuver has no noticeable impact on the uncertainty reduction. The improvement resulting from the control action strongly depends on the uncertainty level at the maneuvering time, making it essential to account for the realistic filter behavior when assessing performance. The low effect of the maneuver for some cases may be attributed to the already low covariance level achieved through the information accumulated from the processed measurements within the navigation filter. This observation reinforces the rationale underlying the information-driven guidance policy. Specifically, the application of a maneuver to further reduce estimation uncertainty is beneficial primarily when the current uncertainty level is sufficiently high to justify the control effort. If the filter alone is capable of achieving satisfactory performance, then an additional control action may be unnecessary and offer limited value.

Finally, in Table 5 one can notice the lower average fuel consumption produced by the POP solutions, which is aligned to the lower performance shown in the trend of the relative range uncertainty.

A concluding remark is that, although in very few cases, the NLP has failed to converge for certain problem initializations, whereas the POP consistently converges to a solution, albeit one that may be sub-optimal. This shows a more robust nature of the POP algorithm in dealing with this maneuver design problem.

INTEGRATION WITHIN A GNC ARCHITECTURE

From an operational standpoint, the findings highlighted in the long-term scenario have significant implications. By leveraging the enhanced information content induced by a specifically planned optimal maneuver, it becomes possible to reduce the frequency of navigation updates without compromising estimation quality. This can lead to several practical advantages, such as reduced onboard computational load, lower demands on ground support and communication bandwidth, and decreased reliance on continuous measurement acquisition or processing. Additionally, fewer navigation windows imply fewer required camera activations or measurement opportunities, which can be beneficial for power-limited missions or when observational opportunities are constrained.

According to this concepts, for future developments, the information-driven guidance policy might be seen as a building block to be integrated in a more complete GNC architecture. The basic idea would be an algorithm aiming at optimizing the planning of the navigation strategy with the final goal of reducing the overall costs. The purpose of the algorithm would then be to decide whether to exploit already planned navigation windows and whether to use a control action to improve estimation performance. The main idea is to consider a nominal mission timeline in which station-keeping maneuvers and relative navigation windows are already scheduled and to let the algorithm plan a strategy to optimize a general cost associated to navigation operations and risks.

In essence, an optimal planning for navigation operations would maintain high estimation performance while reducing navigation activities. This supports more autonomous and resource-efficient mission operations, particularly valuable for deep-space or cislunar scenarios, where spacecraft capabilities might be limited and robustness is essential.

When dealing with a space-based measurement acquisition process, it is crucial to consider that rapid changes in the relative geometry between the observer and the target can lead to periods in which acquiring useful information becomes unfeasible. This limitation reinforces the importance of maximizing estimation performance during the time windows in which measurements are actually available. Efficiently exploiting these opportunities through optimized guidance policies might become critical to ensure navigation accuracy.

CONCLUSIONS

The goal of this work was to evaluate the potential and limitations of an information-driven guidance scheme in the context of inter-satellite angles-only navigation in a three-body dynamics. The analysis and results presented highlighted how guidance policies that explicitly account for state estimation performance can outperform traditional station-keeping approaches when properly integrated with mission constraints such as fuel consumption and reference tracking. Various test cases were considered, including halo and NRHO configurations, with different levels of dynamical complexity and constraint activation, allowing a comprehensive assessment of the proposed strategies.

A key finding of this work is that optimal impulsive maneuvers, computed both through fully nonlinear or polynomial-based optimization frameworks, can significantly reduce the relative range uncertainty even with a single action, and in some cases maintain improved performance over extended periods. This advantage is evident even when realistic errors are introduced in orbit determination and maneuver execution, confirming the robustness of the proposed methodology.

Importantly, the results confirmed that even a single well-designed maneuver can produce long-lasting improvements in estimation performance, potentially allowing operators to reduce the frequency of navigation updates without compromising accuracy. This insight offers clear operational benefits, especially in missions where measurement acquisition opportunities are limited or where resource usage must be minimized.

This study has demonstrated the feasibility and effectiveness of a guidance and navigation co-design process aiming to enhance information acquisition in angles-only navigation. The formulation and application of observability-aware cost functions within constrained optimization problems represent a step forward in enhancing autonomy and robustness in future space missions. The proposed framework lays the groundwork for further research on adaptive, information-driven trajectory planning under realistic mission constraints.

REFERENCES

[1] J. A. Greaves and D. J. Scheeres, "Autonomous Optical-Only Spacecraft-to-Spacecraft Absolute Tracking and Maneuver Classification in Cislunar Space," *Journal of Guidance, Control, and Dynamics*, Vol. 46, No. 11, 2023, pp. 2092–2109.

[2] D. C. Woffinden and D. K. Geller, "Observability criteria for angles-only navigation," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 45, No. 3, 2009, pp. 1194–1208.

[3] J. A. Greaves and D. J. Scheeres, "Autonomous Information Gathering for Distributed Space Systems Using Relative Optical Sensing," *Journal of Guidance, Control, and Dynamics*, 2025, pp. 1–11.

[4] V. Szebehely and E. Grebenikov, "Theory of Orbits-The Restricted Problem of Three Bodies.," *Soviet Astronomy, Vol. 13, p. 364*, Vol. 13, 1969, p. 364.

[5] M. Berz, *Modern Map Methods in Particle Beam Physics*. Academic Press, 1999.

[6] J. B. Lasserre, "Global Optimization with Polynomials and the Problem of Moments," *SIAM Journal on Optimization*, Vol. 11, Jan. 2001, pp. 796–817, 10.1137/S1052623400366802.

[7] J. B. Lasserre, *Moments, Positive Polynomials And Their Applications*, Vol. 1 of *Imperial College Press Optimization Series*. London, UK: Imperial College Press, 1 ed., 2009.

[8] J. L. Junkins and P. Singla, "How nonlinear is it? A tutorial on nonlinearity of orbit and attitude dynamics," *The Journal of the Astronautical Sciences*, Vol. 52, 2004, pp. 7–60.