SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Executive Summary of the Thesis

Safe Guidance Under Uncertainty via Stochastic Optimal Control for Autonomous Small-Body Proximity Operations

Laurea Magistrale in Space Engineering - Ingegneria Spaziale

Author: Luca Frassinella

Advisor: Prof. Francesco Topputo

Co-advisor: Prof. Kenshiro Oguri

Academic year: 2024-2025

1. Introduction

Close-proximity operations around small bodies are influenced by highly nonlinear dynamics and uncertainties that can perturb the spacecraft's nominal trajectory. Autonomy and robustness in these regimes are especially crucial due to the long ground communication delays and the fast dynamics.

This thesis designs safe and autonomous science-orbit transfers around small bodies by applying stochastic optimal control, integrating linear covariance steering within a solution framework based on sequential convex programming (SCP). Robust trajectories and feedback control policies are simultaneously designed to probabilistically guarantee safety under navigation uncertainty, maneuver execution errors and unmodeled accelerations. Moreover, safety under temporary loss of control authority is explicitly addressed. Off-nominal ballistic trajectories following a missed-thrust event are constrained to remain outside a predefined safety region with high probability.

Two stochastic optimization strategies are compared. The first minimizes the statistical fuel cost through the ΔV_{99} metric. The second introduces a convex bound on nonlinear dynamical effects using higher-order dynamical information with State Transition Tensors (STTs) and semi-analytical measures of nonlinearity, with the objective of reducing the mismatch between linearized covariance predictions and the actual nonlinear behavior.

The proposed framework is demonstrated with two representative Apophis transfers between a Periodic (PTO) and a Resonant Terminator Orbit (RTO), and validated via nonlinear Monte Carlo simulations.

2. Methodology

This section summarizes the methodological framework adopted in this thesis. The core architecture is illustrated in Figure 1, which outlines the safe autonomy guidance and control framework adopted in this work [4]. A chance-constrained optimal control problem is solved on the ground using sequential convex programming, incorporating models of the spacecraft dynamics, uncertainties, and mission constraints to compute safe nominal trajectories and associated linear feedback control policies. The resulting policies are validated via nonlinear Monte Carlo simulations and subsequently uploaded onboard. During flight, the spacecraft estimates its state through the onboard navigation filter

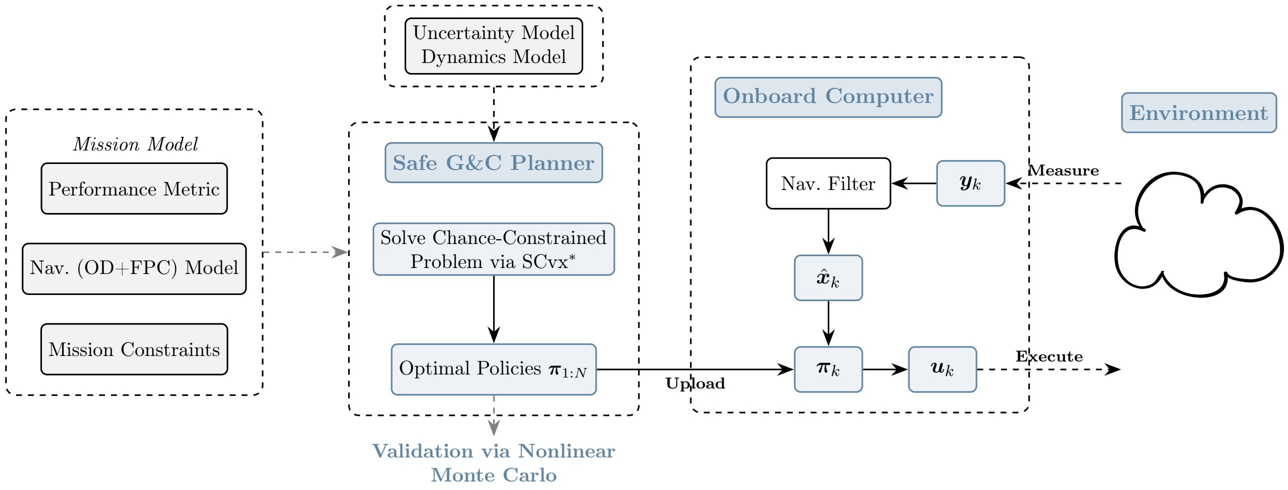

Figure 1: Safe Autonomy framework for chance-constrained Guidance and Control [4].

and applies impulsive control actions to safely execute the transfer.
To initialize the stochastic optimization, a deterministic reference trajectory is first generated using an Adaptive-Mesh sequential convex programming framework [2], which directly optimizes both impulsive maneuvers and the time discretization. This approach places the control actions at optimal epochs, approaching as closely as possible the theoretical optimum provided by Pontryagin's Minimum Principle while retaining numerical efficiency. The method enables the generation of free-phase optimal trajectories from the departing PTO to the arrival 6:1 RTO with bounded time of flight, optimizing the departure and arrival locations as well as the transfer duration.

2.1. Stochastic Optimal Problem Formulation

This thesis formulates the science-orbit transfer under uncertainty as a stochastic optimal control problem, incorporating uncertainty arising from i) initial Gaussian state dispersion , ii) maneuver execution errors with the Gates model, iii) navigation (OD) uncertainty, iv) unmodeled accelerations modeled as a Brownian motion.
The spacecraft dynamics are expressed as a set of nonlinear stochastic differential equations (SDEs):

$$\mathrm{d}\boldsymbol{x} = \left[\boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}, t) + B\tilde{\boldsymbol{u}}\,\right]\mathrm{d}t + G(\boldsymbol{x})\mathrm{d}\boldsymbol{w}(t)$$

$$\boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}, t) = \boldsymbol{f}_0(\boldsymbol{x}, t) + B\boldsymbol{u}, \quad B = \left[0_{n_u \times n_u}^\top, I_{n_u}^\top\,\right]^\top$$

where $\boldsymbol{f}_0(\cdot)$ denotes the natural dynamics modeled with the ANH3BP [7], $\mathrm{d}\boldsymbol{w}(t) \in \mathbb{R}^{n_w}$ is a standard Brownian motion, and $G(\boldsymbol{x})$ maps the Brownian motion to the dynamics of the system.
The objective of the stochastic optimal control problem is to minimize a statistical performance metric while satisfying a set of statistical constraints. This thesis considers path chance constraints both in discrete and continuous-time formulations.

Linear Covariance Steering Formulation

To solve the originally intractable stochastic problem via sequential convex programming, the nonlinear stochastic dynamics are linearized about the reference trajectory at each SCP iteration. Under the assumption of locally linear dynamics and Gaussian-distributed uncertainties, the evolution of the state distribution can be characterized entirely by its mean and covariance. Following this approximation, linear covariance steering is utilized to control both the nominal trajectory and the state distribution, while minimizing the statistical cost. The control policy is modeled as affine in the estimated state,

$$\boldsymbol{u}_k = \bar{\boldsymbol{u}}_k + K_k\left(\hat{\boldsymbol{x}}_k - \bar{\boldsymbol{x}}_k\right),$$

where $\bar{\boldsymbol{u}}_k$ and $\bar{\boldsymbol{x}}_k$ denote the nominal control and state sequences, K_k are linear feedback gains, and $\hat{\boldsymbol{x}}_k$ is the state estimate provided by a Kalman filter.
The decision variables of each convex subproblem are the nominal state and control sequences,

$$\bar{\boldsymbol{X}} = \left[\bar{\boldsymbol{x}}_0^\top \quad \dots \quad \bar{\boldsymbol{x}}_N^\top\right]^\top, \quad \bar{\boldsymbol{U}} = \left[\bar{\boldsymbol{u}}_0^\top \quad \dots \quad \bar{\boldsymbol{u}}_{N-1}^\top\right]^\top,$$

together with the feedback gains $\{K_k\}_{k=0}^{N-1}$.
This work uses a block-Cholesky formulation to propagate the filtered state dynamics and express covariance constraints in a convex form, enabling efficient solution via SCP [5].
Path constraints are imposed as discrete-time chance constraints on the control magnitude, while terminal conditions are enforced on the final distribution in terms of mean and covariance,

$$\mathbb{E}[\boldsymbol{x}_N] = \boldsymbol{x}_f, \qquad \mathrm{Cov}[\boldsymbol{x}_N] \preceq P_f,$$

where $\boldsymbol{x}_f$ and P_f denote the desired terminal mean and covariance.

2.2. Continuous-Time Passive Safety Chance Constraint

This section defines the methodology adopted to enforce passive safety guarantees in the event of a loss of control authority during the autonomous science-orbit transfer, conceptually illustrated in Figure 2.

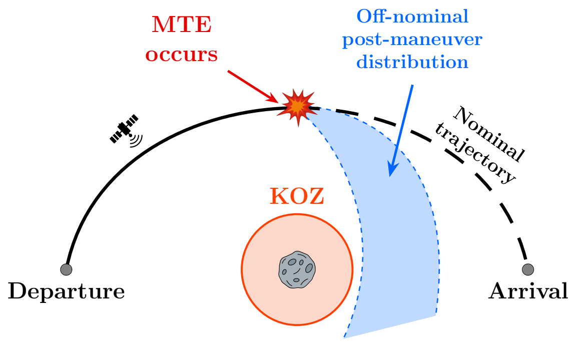

Figure 2: Conceptual illustration of the proposed safety framework.

Following a missed-thrust event (MTE) at each discretization node k, the corresponding post-maneuver distribution is propagated ballistically over a safety horizon of 7 days and required to remain outside a predefined keep-out zone (KOZ) around Apophis.
To prevent inter-sample violations and ensure continuous-time satisfaction of the safety requirement, a convex continuous-time chance constraint is formulated following [1]. An instantaneous penalty function is defined as

$$\Lambda(t) \triangleq \max\Big(0,\ r_{\mathrm{safe}} - \Big[\|\bar{\boldsymbol{r}}(t)\|_2 - \sqrt{m_{\chi^2}(\epsilon_{\mathrm{safe}},3)}\,\|P_r^{1/2}(t)\|_2\Big]\Big)^n$$

where $n > 1$ and $m_{\chi^2}(\epsilon_{\mathrm{safe}}, 3)$ denotes the chi-square quantile with three degrees of freedom evaluated at probability $1-\epsilon_{\mathrm{safe}}$.
The state is augmented with an auxiliary scalar variable $q(t)$ accumulating safety violations,

$$\dot{q}(t) = \Lambda(t), \qquad q(t_k) = 0.$$

Since $\Lambda(t) \geq 0$ by construction, enforcing

$$q(t_k + t_{\mathrm{safety}}) = 0, \quad \forall k,$$

is equivalent to requiring $\int_{t_k}^{t_k+t_{\mathrm{safety}}} \Lambda(t)\,\mathrm{d}t = 0$, which guarantees continuous-time satisfaction of the passive safety constraint over the entire safety horizon.

2.3. Cost Functions

Two different performance metrics are considered. The first minimizes the statistical fuel cost in terms of the 99%-quantile of the control effort by minimizing the ΔV_{99} metric,

$$J^{(\min \Delta V)} = \sum_{k=0}^{N-1} \mathcal{Q}_{\|\boldsymbol{u}_k\|_2}(p = 0.99). \tag{1}$$

The second formulation employs the cost proposed in [6], which minimizes higher-order nonlinear errors neglected by the first-order linearization,

$$J^{(\min \mathrm{NL})} = \max_k \left[(1-\lambda)\,\tilde{\epsilon}_{r,k} + \lambda\,\tilde{\epsilon}_{v,k}\right], \tag{2}$$

where $\tilde{\epsilon}_{r,k}$ and $\tilde{\epsilon}_{v,k}$ are convex bounds of semi-analytical measures of nonlinearity computed from State Transition Tensors (STTs) to approximate the second-order remainders of the nonlinear dynamics, and $\lambda = 0.5$ is a scalar that balances position and velocity contributions, attributing the same weight in this work. This formulation aims at minimizing the discrepancy between linear covariance predictions and the true nonlinear evolution of the state distribution.

2.4. Solution Method via Sequential Convex Programming

Once the originally non-convex problem is reformulated in convex block-Cholesky form [5], it is solved using Sequential Convex Programming. In particular, this work uses the SCvx* algorithm [3], which embeds successive convexification within an Augmented Lagrangian framework. Lagrange multipliers and a quadratic

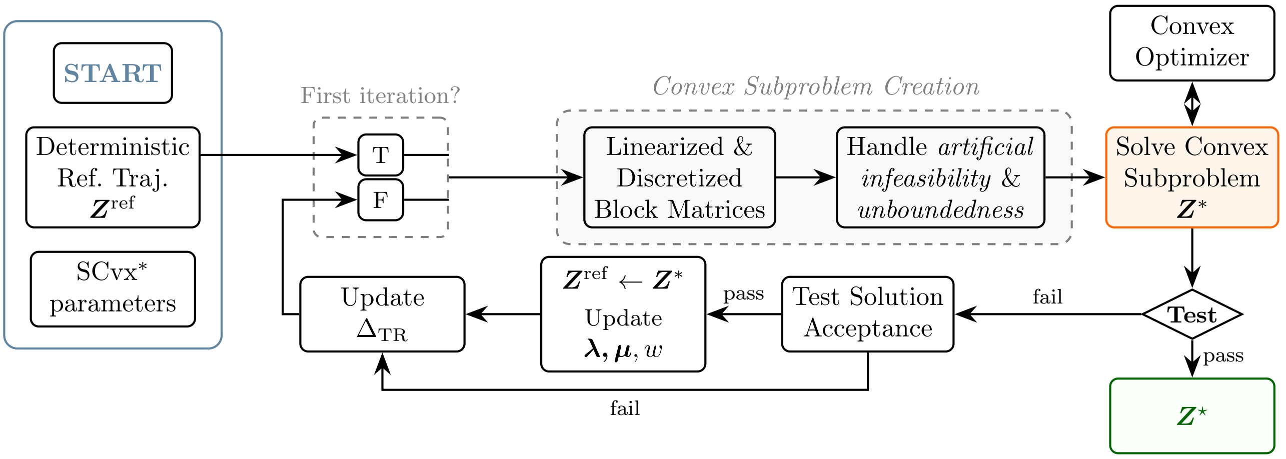

Figure 3: Block-Cholesky Sequential Covariance Steering framework via SCvx* [3, 5].

penalty are introduced for the relaxed constraints, providing theoretical guarantees of convergence to a feasible local solution of the original non-convex problem.
Starting from the deterministic reference trajectory, the stochastic problem is linearized and discretized at each iteration. To prevent *artificial unboundedness* caused by linearization, trust-region constraints are imposed on the decision variables. To address artificial infeasibility, slack variables are introduced to relax selected constraints. The cost function is augmented with penalty terms weighted by Lagrange multipliers, which penalizes constraint violations.
The resulting finite-dimensional convex subproblem is then solved to update the reference trajectory. The process is repeated until convergence in both feasibility and optimality is achieved. Figure 3 illustrates the overall iterative structure of SCvx*.
At each iteration, the convex subproblem to be solved is in the form:

$$\begin{aligned} \min_{\boldsymbol{Z},\boldsymbol{\xi}} \quad & J_{aug}(\boldsymbol{Z},\boldsymbol{\xi}) = J(\boldsymbol{Z}) + J_{pen}(\boldsymbol{\lambda},\boldsymbol{\xi}) \\ \text{s.t.} \quad & \boldsymbol{g}_{\mathrm{eq,affine}} = 0, \quad \tilde{\boldsymbol{g}}_{\mathrm{eq,relaxed}} = \boldsymbol{\xi} \\ & \boldsymbol{h}_{\mathrm{cvx}} \le 0, \quad \|\boldsymbol{Z} - \boldsymbol{Z}^{\mathrm{ref}}\|_\infty \le \Delta_{\mathrm{TR}}, \end{aligned}$$

Here, $\boldsymbol{g}_{\mathrm{eq,affine}} = 0$ includes the linearized filtered state dynamics and control constraints, $\boldsymbol{h}_{\mathrm{cvx}} \le 0$ represents the convex terminal covariance constraints, and $\tilde{\boldsymbol{g}}_{\mathrm{eq,relaxed}}$ collects the relaxed terminal mean and continuous-time passive safety constraints.

3. Numerical Results

This section presents the numerical results obtained by applying the previously introduced framework to two PTO-RTO transfers around Apophis. The solutions are validated through nonlinear Monte Carlo simulations by propagating the full nonlinear SDEs under the computed control policies.
The first case study involves a multi-revolution transfer. The robust nominal trajectories for both optimization strategies are showcased in Figure 4.

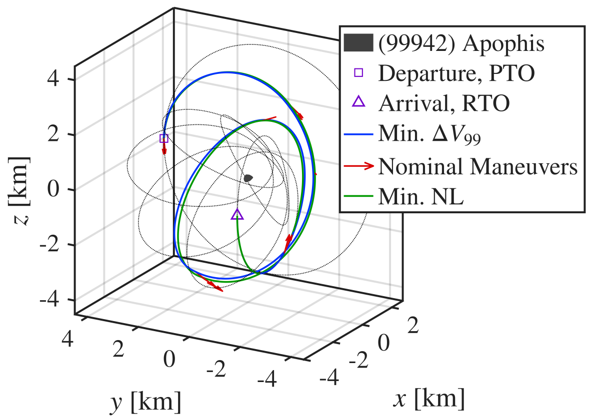

Figure 4: Robust nominal trajectories (Case Study 1).

Although the nominal trajectories are geometrically similar, the two solutions differ significantly in their control structure, particularly in the trajectory correction maneuver (TCM) policies, as shown in Figure 5. The minimum ΔV_{99} formulation primarily relies on nominal feedforward maneuvers, with limited feedback TCMs to reduce the statistical fuel cost. In contrast, the minimum nonlinearity solution introduces more aggressive TCMs to actively control the dispersion and maintain it within the locally linear and Gaussian regime.

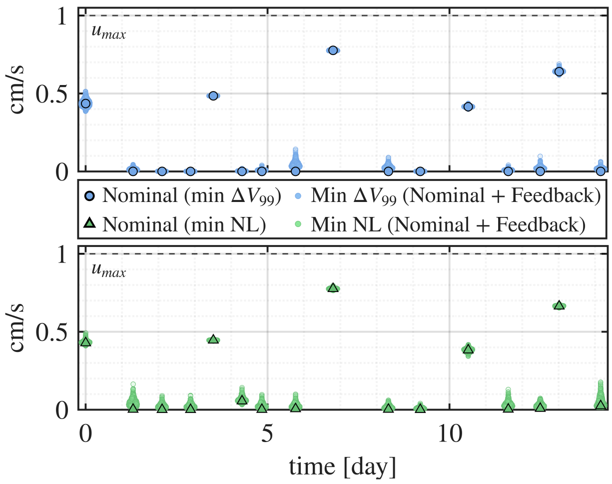

Figure 5: MC control magnitudes at discrete maneuver epochs for Min. ΔV_{99} and Min. NL strategies (Case Study 1).

As a consequence, in the minimum ΔV_{99} case the nonlinear MC samples progressively deviate from the linear prediction, leading to a terminal distribution that does not fully satisfy the target statistics (Figure 6a). This discrepancy is the result of the accumulation of nonlinear effects in the ANH3BP dynamics, and the real distribution quickly becomes non-Gaussian. Conversely, the minimum nonlinearity mitigates these effects, resulting in a much closer agreement between the predicted and sample terminal distributions (Figure 6b).
This improved consistency is achieved at the expense of a slightly higher fuel cost, although its increase remains limited to a few mm/s, as shown in Table 1.

ΔV Statistics [cm/s]

	Min. ΔV_{99}	**Min. NL**
Nominal ΔV	2.75	2.80
Mean MC ΔV	2.91	3.06
MC ΔV_{99}	3.17	3.41

Table 1: Monte Carlo ΔV statistics for minimum ΔV_{99} and minimum nonlinearity objectives (Case Study 1).

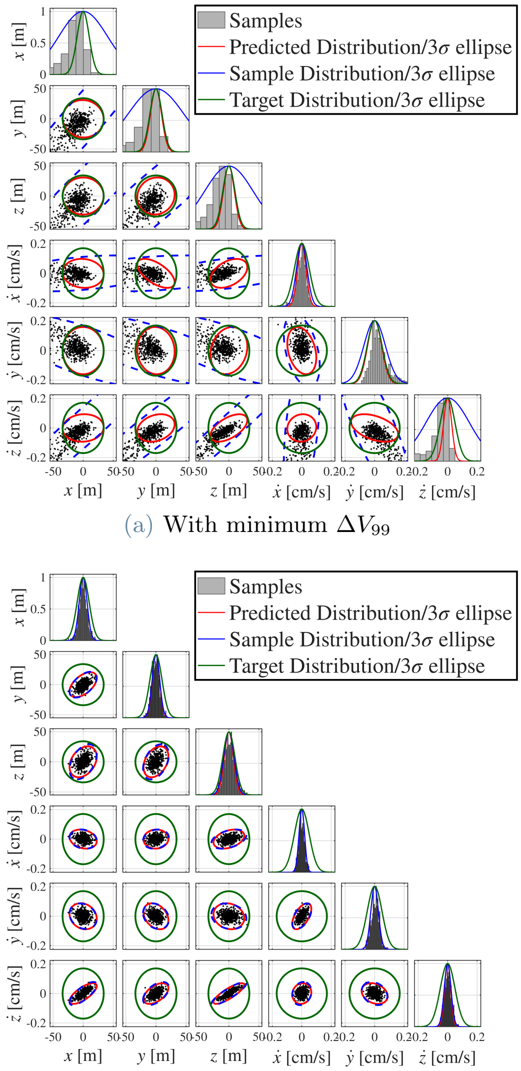

(a) With minimum ΔV_{99}

(b) With minimum nonlinearity

Figure 6: State distribution at target node (Case Study 1).

The second case study considers a flyby transfer with a close approach to Apophis. The robust nominal trajectories are showcased in Figure 8. In this case, the minimum ΔV_{99} strategy results in a higher and more conservative nominal flyby altitude, due to the safety constraints and the larger dispersions associated with its solution. In contrast, the minimum nonlinearity solution allows a closer nominal approach.
The nonlinear Monte Carlo results for the nominal operations exhibit trends similar to Case Study 1 and are omitted for brevity. Figure 7 shows the comparison between the Monte Carlo trajectories obtained in open loop (i.e., with $K_k = 0 \ \forall k$) and in closed loop for the minimum nonlinearity solution, highlighting the essential role of TCMs for robust autonomous

science-orbit transfers under uncertainty. In the open loop case, the dispersion grows significantly along the transfer, leading to large deviations from the nominal trajectory, while the closed loop simulation shows active regulation of the dispersion through the feedback controls, maintaining the distribution within the locally linear regime and ensuring convergence toward the target distribution.

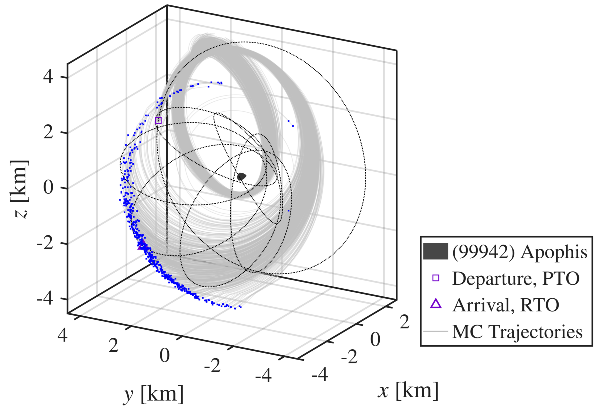

(a) Open Loop (i.e., $K_k = 0\ \forall k$).

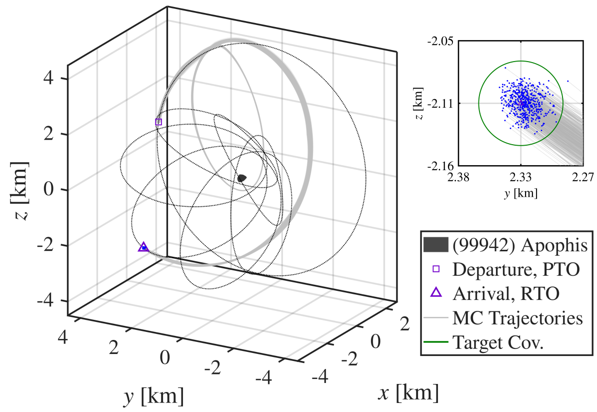

(b) Closed Loop (Min. NL).

Figure 7: Open vs Closed Loop MC Trajectories (Case Study 2).

In this case study, greater insight is obtained by analyzing off-nominal scenarios, where MTEs are simulated at each maneuver epoch. A separate off-nominal Monte Carlo analysis is performed and the resulting ballistic trajectories, propagated from each node over the full safety horizon, are shown in Figure 9. Both strategies are effective in guaranteeing passive safety, with only minimal KOZ violations exhibited by the minimum nonlinearity solution.

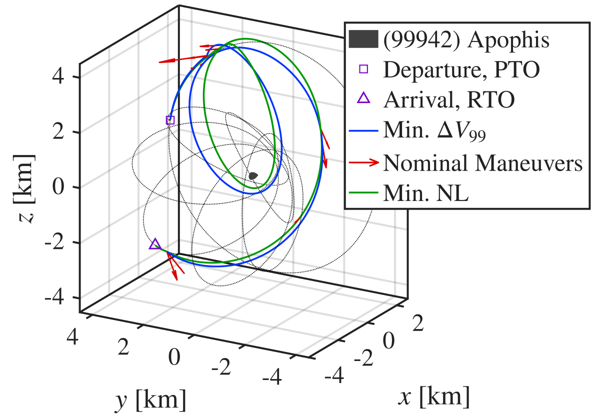

Figure 8: Robust nominal trajectories (Case Study 2).

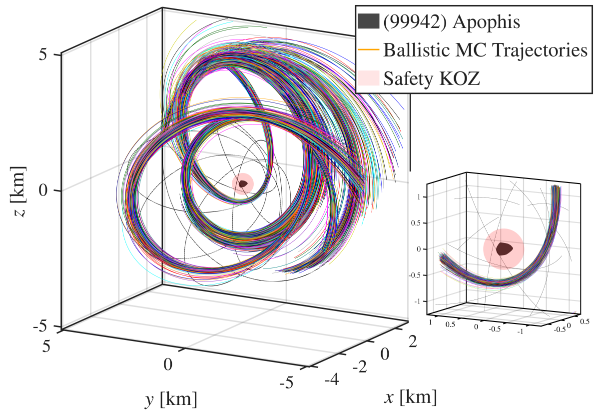

(a) With minimum ΔV_{99}

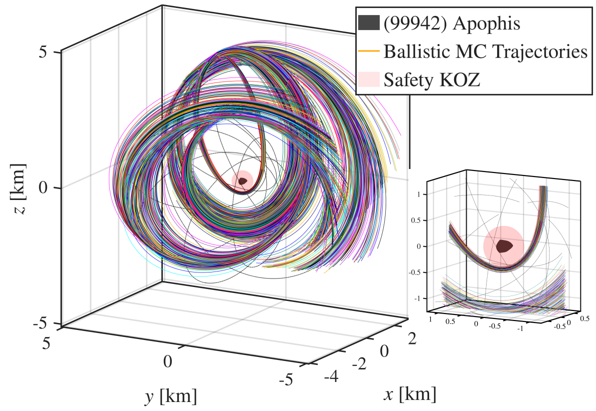

(b) With minimum nonlinearity

Figure 9: Ballistic trajectories after MTEs (Case Study 2).

Although the minimum nonlinearity solution starts from smaller post-maneuver covariances, the resulting ballistic dispersion appears larger in the MTE simulations. This is not due to higher uncertainty levels, but to the structure of the control profile. The minimum ΔV_{99} trajectory contains several nearly ballistic segments with limited corrective action, so MTEs at those nodes have minor impact and the overall deviations remain more contained. This indicates that robustness to missed-thrust events depends not only on covariance magnitude, but also on the frequency of control actions. To this end, introducing sparsity-promoting techniques within the minimum nonlinearity framework to reduce TCMs could provide a more favorable trade-off between terminal distribution robustness and passive safety under MTE scenarios.

4. Conclusions

This thesis presented a stochastic optimal control framework for safe and autonomous proximity operations around small bodies, integrating chance-constrained linear covariance steering within a sequential convex programming solution framework.

Two performance metrics were compared. The minimum ΔV_{99} strategy, which aims at minimizing the statistical fuel cost, resulted in solutions with limited corrective feedback and reduced control effort. However, due to the highly nonlinear regime, the accumulation of nonlinear effects lead to significant discrepancies between the predicted and actual state distributions. In contrast, the minimum nonlinearity formulation, which bounds higher-order dynamical errors by using measures of nonlinearity based on STT information, resulted in improved consistency between linear covariance predictions and nonlinear Monte Carlo samples, at the cost of a marginal increase in fuel consumption.

Furthermore, the thesis introduced a passive safety guarantee in case of missed-thrust events through a continuous-time chance-constrained formulation, demonstrating the capability of both strategies to ensure probabilistic safety even under temporary loss of control authority.

Future work should extend the framework beyond Gaussian assumptions by integrating non-Gaussian distribution steering techniques. Additional directions include the implementation of more computationally efficient stochastic optimization formulations, a more accurate modeling of the navigation and state estimation process and the adoption of higher-fidelity dynamical models to better capture the true small-body environment.

5. Acknowledgements

The author would like to sincerely thank Prof. Francesco Topputo and Prof. Kenshiro Oguri for their continuous guidance, insightful discussions, and support throughout this research.

The author also gratefully acknowledges Politecnico di Milano and Purdue University for providing the academic and research environment in which this work was carried out.

References

[1] Purnanand Elango, Dayou Luo, Abhinav G Kamath, Samet Uzun, Taewan Kim, and Behçet Açıkmeşe. Successive Convexification for Trajectory Optimization with Continuous-Time Constraint Satisfaction. *arXiv preprint arXiv: 2404.16826*, 2024.

[2] Naoya Kumagai and Kenshiro Oguri. Adaptive-Mesh Sequential Convex Programming for Space Trajectory Optimization. *Journal of Guidance, Control, and Dynamics*, 47(10):2213–2220, 2024.

[3] Kenshiro Oguri. Successive Convexification with Feasibility Guarantee via Augmented Lagrangian for Non-Convex Optimal Control Problems. In *2023 62nd IEEE Conference on Decision and Control (CDC)*, pages 3296–3302. IEEE, 2023.

[4] Kenshiro Oguri. Chance-Constrained Control for Safe Spacecraft Autonomy: Convex Programming Approach. In *2024 IEEE American Control Conference*, 2024.

[5] Kenshiro Oguri and Gregory Lantoine. Convex Block-Cholesky Approach to Risk-Constrained Low-thrust Trajectory Design under Operational Uncertainty. *arXiv preprint arXiv: 2602.18416*, 2026.

[6] Daniel C Qi and Kenshiro Oguri. Convex Bound of Nonlinear Dynamical Errors for Stochastic Optimal Control. *arXiv preprint arXiv: 2510.21975*, 2025.

[7] DJ Scheeres and Francesco Marzari. Spacecraft Dynamics in the Vicinity of a Comet. *The Journal of the astronautical sciences*, 50(1):35–52, 2002.