Statistical Trajectory Optimization Under Uncertainty for Icy Moon Exploration with Probability of Impact Constraints

Gabriele Nuccio[(1)(2)], Kenshiro Oguri[(2)], Francesco Topputo[(1)]

[(1)] *Department of Aerospace Science and Technology, Politecnico di Milano*
Milan, Italy
Email: gabriele.nuccio@mail.polimi.it, francesco.topputo@polimi.it

[(2)] *School of Aeronautics and Astronautics, Purdue University*
West Lafayette, Indiana, USA
Email: koguri@purdue.edu

Abstract – **Trajectory design for icy moon exploration in highly perturbed environments such as the Jovian system requires explicit treatment of uncertainties, particularly during repeated close flybys that demand precise targeting and safety guarantees. This work presents a sequential covariance steering approach based on the SCvx* algorithm to design simultaneously optimal transfers and an associated trajectory-correction policy while probabilistically guaranteeing safety constraints under statistical uncertainties. Two convex formulations are implemented and compared within the same framework: a block Cholesky method and a full covariance approach. At each flyby, B-plane geometric targeting constraints are imposed while a safety chance constraint ensures a minimum flyby altitude with high probability even under complete loss of control authority. To reduce the number of active maneuvers, an $\ell_{1,p}$ regularized sparsification algorithm is introduced, inspired by recent developments in hands-off covariance steering. The approach is demonstrated on representative Jupiter–Europa transfers in the Circular Restricted Three-Body Problem (CR3BP) and validated through nonlinear Monte Carlo simulations which confirm the robustness of the transfers when reintroducing the full nonlinear stochastic dynamics.**

I. INTRODUCTION

Among the most compelling targets in the search for habitable environments beyond Earth are the icy moons of the outer Solar System, whose frozen crusts and subsurface oceans suggest conditions potentially favorable to life [1]. In particular, Europa has emerged as a primary scientific focus, motivating dedicated exploration campaigns such as NASA's Europa Clipper mission [2]. Meeting the associated science objectives demands a sequence of close flybys and intricate transfers carried out within a strongly perturbed multi-body environment.

From an operational standpoint, repeated-flyby tours in the Jovian or Saturnian systems are intrinsically challenging. The combination of multi-body gravity, highly nonlinear dynamics driven by frequent close encounters, and hostile environmental conditions (e.g., radiation) makes trajectories extremely sensitive to even small dispersions. In this scenario, maneuver execution and navigation errors can amplify across successive encounters, leading to missed observation opportunities and, in the worst case, to unsafe flyby geometries and potential surface impacts at later encounters. Communication delays and constrained visibility windows further motivate a design philosophy in which robustness and onboard autonomy are necessary for mission success.

Traditional mission design typically starts from deterministic reference trajectories and addresses uncertainty *a posteriori* through heuristic margins and extensive navigation analyses. In contrast, this work incorporates statistical uncertainty directly into the mission design process by combining chance-constrained control with sequential convex programming (SCP) in a tractable covariance steering formulation. Stochastic Optimal Control (SOC) approaches based on SCP have already been applied to nonlinear space mission problems [3–6], naturally leading to covariance steering (CS) formulations in which both the mean trajectory and the state dispersion are explicitly controlled under chance constraints [7].

Despite these advances, the literature still lacks a sequential covariance steering framework specifically tailored to icy moon multi-flyby trajectory design, where repeated close encounters and probability of impact (PoI) requirements must be handled within a unified optimization setting.

This work addresses that gap by developing a chance-constrained sequential covariance steering framework for robust trajectory design in representative Jupiter–Europa transfers. In addition, the framework is extended by promoting sparsity of the control maneuvers using a strategy based on hands-off covariance steering [8]. While this approach has been recently introduced and demonstrated on representative test cases, this paper applies it, to the best of the authors' knowledge, for the first time to a relevant space mission design problem involving nonlinear dynamics.

The paper is organized as follows: Section II introduces the mathematical model. Section III gives an overview of the scenario and the design of the deterministic reference trajectories. Section IV presents the complete formulation of the chance-constrained trajectory optimization problem. Section V reports the numerical results. Finally, Section VI summarizes the main findings and outlines possible future developments.

II. Mathematical Model

This section introduces the mathematical model adopted in this work. First, the B-plane representation used for flyby targeting and probabilistic impact assessment is recalled. Then, it reports the equations of motion and the dynamical models utilized for the selected application.

A. B-Plane Definition

The B-plane provides a compact geometric representation of a hyperbolic flyby and is widely used for encounter targeting, uncertainty mapping, and impact-risk analysis. This representation is especially convenient because the encounter geometry can be described by the components of the B-plane vector along the $\hat{\mathbf{R}}$ and $\hat{\mathbf{T}}$ directions:

$$\boldsymbol{b}(\boldsymbol{x}) = \begin{bmatrix} b_R & b_T \end{bmatrix}^\top . \tag{1}$$

In this paper, B-plane coordinates are used both for flyby targeting and for probabilistic impact requirements. For the latter purpose, it is convenient to work with the scaled B-plane coordinates, thus an impact corresponds to $\|\boldsymbol{b}\|_2 < R_{\text{moon}}$ [9]. The reader is referred to [10] for the complete derivation of the B-plane parameters used in this work.

B. Equations of Motion

Let the spacecraft state vector be denoted by $\boldsymbol{x} \in \mathbb{R}^{n_x}$ and the control vector by $\boldsymbol{u} \in \mathbb{R}^{n_u}$. In the present application, the state is expressed in Cartesian coordinates and collects spacecraft position and velocity. The equations of motion for spacecraft orbital dynamics under control are expressed as a control-affine system:

$$\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}, t) = \boldsymbol{f}_0(\boldsymbol{x}, t) + B\boldsymbol{u}(t), \quad B = \begin{bmatrix} \mathbf{0}_{3\times3} \\ \boldsymbol{I}_3 \end{bmatrix} \tag{2}$$

where $\boldsymbol{f}_0(\cdot)$ denotes the uncontrolled dynamics and B maps the control acceleration or impulse into the velocity components of the state.

In this work, the control action is modeled as a finite sequence of impulsive maneuvers. In particular, $\boldsymbol{u}_k \equiv \Delta\boldsymbol{v}_k$ denotes the impulsive velocity increment applied at the prescribed maneuver epoch t_k, with $k \in \mathbb{Z}_0^{N-1}$:

$$\boldsymbol{u}(t) = \sum_{k=0}^{N-1} \boldsymbol{u}_k \, \delta(t - t_k) \tag{3}$$

where $\delta(\cdot)$ denotes the Dirac delta function and N is the planning horizon of the control input opportunities. The impulsive control model is naturally aligned with the selected application [11].

C. Dynamical Model

The Jupiter–Europa Circular Restricted Three-Body Problem (CR3BP) is adopted as the dynamics model to capture the key multi-body effects driving same-body transfers. The equations of motion are written in nondimensional form in a rotating frame, using the standard CR3BP scaling, where the uncontrolled dynamics $\boldsymbol{f}_0(\cdot)$ is given by [12]:

$$\boldsymbol{f}_0(\boldsymbol{x}) = \begin{bmatrix} \boldsymbol{v} \\ \boldsymbol{f}_a(\boldsymbol{x}) \end{bmatrix}, \quad \boldsymbol{f}_a(\boldsymbol{x}) = \begin{bmatrix} 2\dot{y} + \dfrac{\partial\Omega}{\partial x} \\ -2\dot{x} + \dfrac{\partial\Omega}{\partial y} \\ \dfrac{\partial\Omega}{\partial z} \end{bmatrix}. \tag{4}$$

The pseudo-potential function Ω is

$$\Omega(x, y, z) = \frac{1}{2}(x^2 + y^2) + \frac{1-\mu}{r_1} + \frac{\mu}{r_2}, \tag{5}$$

where

$$r_1 = \sqrt{(x+\mu)^2 + y^2 + z^2}, \tag{6a}$$

$$r_2 = \sqrt{(x-1+\mu)^2 + y^2 + z^2}, \tag{6b}$$

$$\mu = \frac{m_2}{m_1 + m_2} \tag{6c}$$

and m_1 and m_2 are the Jupiter and Europa masses, respectively.

III. Scenario and Reference Trajectories

This section introduces the mission scenario considered in this work and summarizes the deterministic reference trajectories used to support the chance-constrained SCP framework developed in the subsequent section.

The considered case study is representative of the science phase of NASA's Europa Clipper mission, which relies on a sequence of close Europa flybys with highly elliptical Jupiter-centered orbits. In this paper, the analysis is restricted to a subset of the Europa science campaign of the 21F31 baseline tour [13], focusing on *ballistic same-body transfers*. Specifically, two representative phases are considered: (i) non-resonant arcs within the *petal rotation* segment, where alternating 1:4$^+$ and 1:5$^-$ pseudo-resonant legs rotate the line of apsides while minimizing the accumulated radiation; and (ii) resonant arcs within the *crank-over-the-top* (COT-1) segment, consisting of 1:6 resonant transfers. The orbital and physical parameters of the Jupiter-Europa system adopted for the considered scenario are summarized in Table 1.

A. Deterministic Reference Generation in the CR3BP

Two deterministic reference trajectories are constructed, a non-resonant one and a resonant one. In both cases, the reference spans two complete petals: it starts immediately after a Europa flyby, evolves through the two sub-

Table 1. Orbital and physical parameters

Jupiter - Europa CR3BP parameters	
Length Unit [km]	671051.2
Velocity Unit [km/s]	13.74
Time Unit [s]	48839.2
Europa gravitational parameter [km^3/s^2]	3200.99
Jupiter gravitational parameter [km^3/s^2]	1.2669×10^8
Mass ratio of Jupiter–Europa system [-]	2.5266×10^{-5}
Europa radius [km]	1564
Jupiter radius [km]	69911

sequent close encounters, and terminates at the following apocenter. The generation procedure follows the approach described in [14] and consists of three main steps:

1. **Patched-conic initialization:** compute an initial guess with the patched-conic approximation consistent with the desired phasing (non-resonant) or resonance (resonant) conditions.
2. **Mapping to the CR3BP:** convert the flyby conditions into rotating-frame initial states and propagate the CR3BP dynamics to obtain a nearly periodic arc.
3. **Differential correction and continuation:** enforce periodicity matching conditions and a prescribed Jacobi constant via a forward-backward shooting differential correction approach. Then, by applying a simple numerical continuation strategy, the solution converges to an operationally meaningful closest-approach altitude.

The non resonant transfer has a total time of flight of $TOF = 39.4$ days while the resonant trajectory presents a total orbital duration of $TOF = 53.4$ days. The resulting deterministic references provide dynamically consistent nominal trajectories that preserve the targeted flyby geometry. They serve as the baselines for dynamics linearization, uncertainty propagation, and probabilistic constraints enforcement in the subsequent chance-constrained SCP framework.

IV. CHANCE-CONSTRAINED SEQUENTIAL COVARIANCE STEERING

This section presents the chance-constrained sequential covariance steering framework adopted to design robust multi-flyby transfers under uncertainty. A chance-constrained control approach is selected as it allows handling unbounded distributions while enforcing safety and operational requirements with a prescribed confidence level [15]. An overview of the proposed architecture is shown in Fig. 1.
The method explicitly incorporates both the force and the uncertainty models and computes statistically optimal nominal trajectories by directly embedding chance constraints into the formulation, while minimizing a statistical objective. Both nominal trajectory and flight-path control (FPC) policies $\pi_{0:N-1}$ are optimized through a sequential convex programming (SCP) algorithm, which iteratively solves tractable convex subproblems [16].
On board, the feedback control policy π_k is uploaded to the flight computer, where a Kalman filter estimates the state $\hat{\boldsymbol{x}}_k$ from the measurements $\boldsymbol{y}_k$, and produces the control action $\boldsymbol{u}_k$ required to execute the safe transfer [6]. The final solution is assessed through nonlinear Monte Carlo simulations, where the full nonlinear stochastic dynamics and uncertainty sources are reintroduced.
The section first outlines the original stochastic optimal control problem. It then introduces the tractable covariance steering reformulation obtained by parameterizing the FPC policies and propagating uncertainty through linearized covariance dynamics.

A. Original Stochastic Optimal Control Problem

The original stochastic optimal control problem is first presented in its original nonlinear form.

Uncertainty Models. The stochastic formulation accounts for the main uncertainty sources encountered in deep-space missions. In this work, uncertainty is modeled through: (i) initial state dispersion due to orbital insertion errors, (ii) maneuver execution errors modeled through the Gates model [17], (iii) stochastic accelerations capturing dynamical mis-modeling and unmodeled perturbations, and (iv) orbit determination (OD) or navigation uncertainty.
Collecting the previous uncertainty sources, the spacecraft dynamics is modeled as a set of stochastic differential equations (SDEs):

$$\mathrm{d}\boldsymbol{x} = [\boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}, t) + B\,\tilde{\boldsymbol{u}}(t)]\,\mathrm{d}t + G(\boldsymbol{x})\,\mathrm{d}\boldsymbol{w}(t), \tag{7}$$

where $\boldsymbol{f}(\boldsymbol{x}, \boldsymbol{u}, t) = \boldsymbol{f}_0(\boldsymbol{x}, t) + B\boldsymbol{u}$. $\boldsymbol{f}_0(\cdot)$ denotes the natural dynamics modeled with the CR3BP, $\tilde{\boldsymbol{u}}(t)$ represents the maneuver execution error mapped into the state dynamics through the input matrix B. Stochastic acceleration due to dynamic mis-modeling can be naturally modeled by a Brownian motion: $G(\boldsymbol{x})\,\mathrm{d}\boldsymbol{w}(t)$, where $G(\cdot)$ is the intensity of disturbances and $\mathrm{d}\boldsymbol{w}(t)$ is a standard Brownian motion vector. The Brownian motion is approximated by discrete-time white noise $\boldsymbol{w}_k$ with a zero-order hold over a small time step $t \in [t_k, t_k + \Delta t_k)$ as [18]:

$$G(\boldsymbol{x})\,d\boldsymbol{w}(t) = G_k(\boldsymbol{x})\frac{1}{\sqrt{\Delta t_{\mathrm{WN}}}}\boldsymbol{w}_k, \tag{8}$$

where Δt_{WN} denotes the time discretization step for approximation of SDE with white noise and G_k is the discretized stochastic disturbances matrix.

Statistical objective function. A standard objective in trajectory optimization is the total control effort. In the

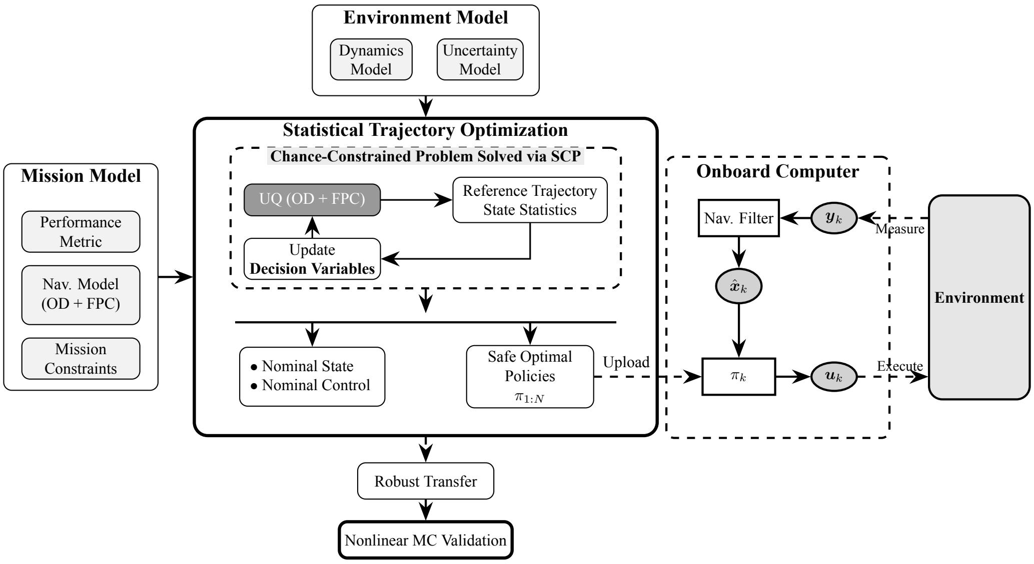

Fig. 1: *Statistical trajectory optimization under uncertainty scheme*

stochastic setting, however, the closed-loop control is itself a random variable because of the feedback policy. Therefore, this work minimizes the 99%-quantile of the closed-loop impulsive control effort,

$$J = \sum_{k=0}^{N-1} Q_{\|\boldsymbol{u}_k\|_2}(p = 0.99), \tag{9}$$

where $Q_X(p)$ denotes the quantile of the random variable X evaluated at probability p. This metric is commonly referred to as ΔV_{99} in space mission design.

Constraints. The constraints considered in this study include both path and terminal requirements. Since state and control are subject to uncertainty, constraints are treated stochastically through chance and distributional formulations. The following constraints are implemented within the considered scenario:

1) *Control chance constraint:* the maneuver magnitude is limited with 99.9 % confidence:

$$\mathbb{P}\left[\|\boldsymbol{u}_k\| \leq u_{\text{max}}\right] \geq 1-\epsilon_u, \qquad k = 0, \cdots, N-1 \tag{10}$$

2) *Terminal distributional constraint:* the final state distribution is shaped by constraining its mean and covariance:

$$\mathbb{E}\left[\boldsymbol{x}(t_f)\right] = \bar{\boldsymbol{x}}_f, \quad \text{Cov}[\boldsymbol{x}(t_f)] \preceq P_f \tag{11}$$

3) *Maximum covariance constraint:* to limit dispersion growth and preserve the validity of linearized covariance propagation, a bound on the covariance is imposed, including at the closest-approach instants:

$$\text{Cov}[\boldsymbol{x}_k] \preceq P_{\text{max}}, \quad k = 0, \cdots, N \tag{12}$$

4) *B-plane and TCA targeting constraint*: rather than enforcing a full-state matching at each flyby, targeting is imposed in the B-plane. This solution reduces the nonlinearities given the linearized mapping [10]. In this work, the mean trajectory is constrained to match the target B-plane coordinates and the time of closest approach (TCA) both extracted from the reference transfer.

$$\boldsymbol{b}(\bar{\boldsymbol{x}}_{FB}) = \bar{\boldsymbol{b}}_{\text{target}}, \quad t_{\text{CA}}(\bar{\boldsymbol{x}}_{FB}) = \bar{t}_{\text{CA, target}} \tag{13}$$

5) *B-plane safety chance constraint*: to enforce planetary protection compliance, a safety chance constraint is introduced to guarantee impact avoidance even under complete loss of maneuverability. For a selected set of nodes along the second half of each petal, i.e., after the apocenter until the subsequent flyby, the trajectory is propagated ballistically from the assumed loss of control instant to the closest approach, and a probabilistic lower bound is imposed on the B-plane miss distance:

$$\mathbb{P}\left[\|\boldsymbol{b}_j\|_2 \geq \rho_{\text{safe}}\right] \geq 1 - \varepsilon_{\text{im}} \tag{14}$$

where $\boldsymbol{b}_j$ denotes the B-plane vector for the j-th final maneuver before the loss of maneuverability and ρ_{safe} is the minimum safe miss distance (set to the Europa radius) and ε_{im} is the allowable impact probability which is selected as 0.1% [19]. This constraint ensures that, with probability at least $1 - \varepsilon_{\text{im}}$, the spacecraft's closest approach distance remains above the prescribed safety threshold even without corrective maneuvers.

The original formulation of the developed chance-constrained optimal control problem can be described as follows:

Problem 1 (Original Chance-Constrained Control Problem**).** *Find the nominal trajectory $\boldsymbol{x}^*(t)$, $\boldsymbol{u}^*(t)$ and the control policies $\pi_{0:N-1}$, that minimize the cost function*

(9) subject to the dynamical constraint (7), path chance constraints (10, 14), maximum covariance constraint (12), terminal constraint (11) and targeting constraint (13) under a sequence of FPC policies informed by the state estimates in Eq. (12) of [6].

B. Tractable Formulation via Covariance Steering

The original problem is generally intractable, as rigorous uncertainty quantification through nonlinear dynamics would be computationally prohibitive. To solve the original SOC problem via sequential convex programming, the dynamics is linearized about the current reference trajectory at each SCP iteration. Under the assumption of Gaussian-distributed uncertainties and locally linear dynamics, the evolution of the state distribution is represented through its first two moments, namely mean and covariance [6].

Therefore, the chance-constrained stochastic optimal control problem is reformulated via *sequential linear covariance steering* to control both the nominal trajectory and the state dispersion, while minimizing the statistical cost metric and satisfying the mission constraints.

In this work, an affine version of the control policy is considered :

$$\boldsymbol{u}_k = \bar{\boldsymbol{u}}_k + K_k(\hat{\boldsymbol{x}}_k - \bar{\boldsymbol{x}}_k) \tag{15}$$

where $\bar{\boldsymbol{x}}_k$ and $\bar{\boldsymbol{u}}_k$ denote the nominal state and control, $\hat{\boldsymbol{x}}_k$ is the state estimate provided by a Kalman filter and K_k is a feedback gain matrix.

C. Covariance Steering Formulations

Two covariance steering formulations are implemented and compared in this work, as this direct comparison has not been thoroughly addressed in the existing literature. Both rely on the affine feedback policy introduced above, but differ in the way the state covariance evolution is parameterized and convexified.

The first one is the *block Cholesky* (BC) formulation. This approach leverages a block lower-triangular Cholesky factor of the covariance matrix, obtained by shaping it after stacking the state and the control statistics [16]. In compact form, the stacked covariance can be written as:

$$P_{\boldsymbol{X}} = P_{\boldsymbol{X}}^{1/2}\left(P_{\boldsymbol{X}}^{1/2}\right)^{\top}, \tag{16}$$

where $P_{\boldsymbol{X}}^{1/2}$ denotes the Cholesky factor of the covariance associated with the complete state history. This parameterization makes several chance constraints tractable in convex form, since they can be expressed through affine functions of the Cholesky factors. However, the resulting formulation involves large linear matrix inequalities (LMIs), whose size grows rapidly with the number of discretization nodes becoming computationally expensive [20].

The second one is the *full covariance* (FC) approach. In this case, the covariance matrices are introduced explicitly as decision variables at each node P_k.

The covariance propagation is then handled through a lossless convexification based on auxiliary variables. The procedure for lossless convexification is defined as follows: denote by $\hat{P}_k$ the covariance of the state estimate, introduce the variables:

$$U_k := K_k\hat{P}_k, \qquad Y_k := U_k\hat{P}_k^{-1}U_k^{\top}, \tag{17}$$

and finally impose the LMI

$$\begin{bmatrix} \hat{P}_k & U_k^{\top} \\ U_k & Y_k \end{bmatrix} \succeq 0. \tag{18}$$

After the optimization, one can recover the feedback gain matrix as $K_k = U_k\hat{P}_k^{-1}$.

This procedure yields a formulation that is computationally less demanding than the block Cholesky approach. A further advantage of the full covariance formulation is that it naturally supports feedback on the latest state estimate, whereas block Cholesky implementations typically rely on feedback policies on the state estimate history [21]. The main drawback is that some chance constraints become nonlinear functions of the covariance variables and therefore require additional convexification techniques within the SCP loop.

Rearranging the state and control variables with the two approaches, it is possible to write the probabilistic constraints in Sec. IV.A in a deterministic and convex form. Here, two deterministic and convex upper bounds of the objective function in (9) are presented for the two formulations:

$$J_0^{\mathbf{BC}}(\boldsymbol{z}) = \sum_{k=0}^{N-1}\left[\|\bar{\boldsymbol{u}}_k\|_2 + m_{\chi^2}(\varepsilon, n_u)\|P_{\boldsymbol{u}_k}^{1/2}\|_2\right] \tag{19a}$$

$$J_0^{\mathbf{FC}}(\boldsymbol{z}) = \sum_{k=0}^{N-1}\|\bar{\boldsymbol{u}}_k\| + m_{\chi^2}(\varepsilon, n_u)\sqrt{\lambda_{\max}(Y_k)} + \epsilon_Y\,\mathrm{tr}(Y_k) \tag{19b}$$

where $m_{\chi^2}(\varepsilon, n_u)$ denotes the square root of the quantile function of the chi-squared distribution with n_u degrees of freedom evaluated at probability $1-\varepsilon$. These bounds are minimized in place of the original cost function for tractability purposes. For complete details of the two formulations (BC and FC), the reader is referred to [16, 21], respectively. In this paper, both are implemented for comparison, while the full covariance formulation is adopted as the baseline for the hands-off extension.

*D. Solution Method via SCvx**

After reformulating the original non-convex problem in a convex form as explained above, it is solved using a SCP framework. In particular, this work adopts the SCvx* algorithm [22] which is based on a sequential convexification loop with an outer update of the Lagrange multipliers, effectively following the augmented Lagrangian (AL) method. The SCvx* algorithm is demonstrated to globally converge to a feasible local optimum of the original problem with a linear convergence rate of the Lagrange multipliers, which take care of penalizing the constraints violations. The statistical cost function is therefore augmented with the penalty term weighted by the Lagrange multipliers $\boldsymbol{\lambda}$ and $\boldsymbol{\mu}$ introduced in Eq. (52) of [21] and presented also in [23].

The convex subproblem is then solved at every iteration of the loop. The procedure is repeated until convergence in both feasibility and optimality metrics is achieved.

Denoting by $\boldsymbol{z}$ the collection of decision variables, including nominal states, controls, feedback gains, covariance-related variables, and slack variables, the generic convex subproblem solved at each iteration can be written as:

Problem 2 (Convex subproblem)**.**

$$\begin{aligned}\min_{\boldsymbol{z},\boldsymbol{\xi},\boldsymbol{\zeta}\geq 0}\quad & J_{\text{aug}}(\boldsymbol{z},\boldsymbol{\xi},\boldsymbol{\zeta}) := J_0(\boldsymbol{z}) + J_{\text{pen}}(\boldsymbol{\xi},\boldsymbol{\zeta}) \\ \text{s.t.}\quad & \tilde{\boldsymbol{g}}(\boldsymbol{z}) = \boldsymbol{\xi}, \quad \tilde{\boldsymbol{h}}(\boldsymbol{z}) \leq \boldsymbol{\zeta}, \quad \boldsymbol{\zeta} \geq 0, \\ & \boldsymbol{g}_{\text{affine}}(\boldsymbol{z}) = 0,\ \boldsymbol{h}_{\text{cvx}}(\boldsymbol{z}) \leq 0, \\ & \|\boldsymbol{z} - \bar{\boldsymbol{z}}\|_\infty \leq \Delta^{(k)}.\end{aligned}$$

Here, $\tilde{\boldsymbol{g}}$ and $\tilde{\boldsymbol{h}}$ are the linearized nonconvex equality and inequality constraints, respectively. The trust-region constraint limits the validity region of the local approximation and mitigates artificial unboundedness. To address artificial infeasibility, two slack variables are introduced: $\boldsymbol{\xi}$ relaxes a subset of the equality constraints (dynamics and B-plane targeting), while $\boldsymbol{\zeta}$ relaxes a subset of the inequality constraints (B-plane safety chance constraint). After each accepted step, the reference trajectory, trust-region radius, and Lagrangian multipliers are updated according to the SCvx* procedure.

E. Hands-Off Covariance Steering

In practical mission operations, each correction maneuver introduces costs beyond propellant consumption, including operational overhead and attitude reorientation for thruster alignment. For this reason, it is often desirable to minimize not only the total control effort but also the number of active control interventions, leading to sparse control profiles in which many maneuver opportunities are left unused.

Building on the full covariance steering formulation, this work integrates the iteratively reweighted $\ell_{1,p}$ (IRL1P) framework proposed by Kumagai and Oguri [8] to promote group sparsity in the feedback component of the correction policy. This choice is motivated by the numerical evidence that, in the considered scenario, the feedback term represents the dominant contribution to the overall trajectory correction maneuver (TCM) profile. Practical modifications are introduced to improve the sparsity-promotion performance in the considered setting while retaining compatibility with the SCvx* inner loop used to solve the convex subproblems. A flowchart summarizing the resulting nested procedure, composed of an outer IRL1P loop and an inner SCvx* loop, is provided in Fig. 2. The main modifications are reported in the following.

At the l-th outer iteration, the augmented objective of the full covariance subproblem is modified as

$$J_{\text{aug}} = J_0(\boldsymbol{z}) + J_{\text{pen}}(\boldsymbol{\xi},\boldsymbol{\zeta}) + \psi \sum_{k=0}^{N-1} w_k^{(l)} \|U_k\|_F, \quad (20)$$

where $\psi > 0$ is a regularization parameter controlling the sparsity promotion and U_k is the full covariance auxiliary variable introduced in Section IV.C. Unlike the original formulation, where the regularization is imposed on Y_k, this work regularizes U_k, which was found to improve sparsity in the considered application.

The reweighting coefficients are updated from the recovered feedback gains as

$$w_k^{(l+1)} = \frac{1}{\|K_k^{(l)}\|_F + \varepsilon}, \qquad w_k^{(1)} = 1, \quad (21)$$

where $\varepsilon > 0$ is a small user-defined parameter. This update assigns larger penalties to maneuvers whose feedback gains are already small, progressively driving negligible correction actions toward zero.

Since numerical optimization generally leaves small residual gains rather than exact zeros, a thresholding mechanism is introduced. If the feedback gain at a maneuver opportunity satisfies

$$\|K_k\|_F < \delta_{\text{th}}, \quad (22)$$

then the corresponding correction action is deactivated. Once a maneuver is nullified, it is kept inactive in the subsequent outer iterations, allowing the sparsity pattern to progressively stabilize. Finally, in the hands-off configuration, the B-plane safety chance constraint is enforced on every maneuver opportunity within each petal up to the subsequent flyby, with the exception of the first one, instead of being restricted to the second half of the petal as in the baseline formulation. This modification is motivated by the sparsification process which nullifies several maneuvers. With the resulting sparse profile, applying the safety constraint over the entire petal remains tractable and yields a more realistic design approach.

F. Nonlinear Monte Carlo Simulation

After the optimization, the solution is verified through nonlinear Monte Carlo (MC) simulations, following the methodology originally proposed by Oguri and Lantoine [16]. Due to the nonlinearity of the simulation, an extended Kalman filter (EKF) replaces the Kalman filter

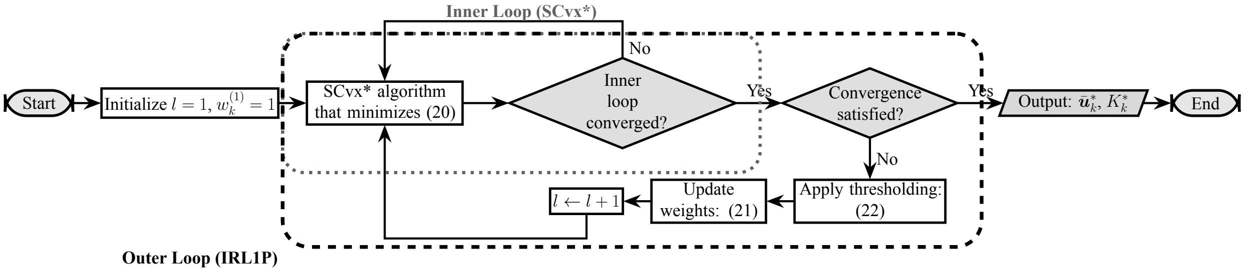

Fig. 2: *Flowchart of the hands-off covariance steering algorithm*

(KF) for the filtering process. Moreover, state propagation is performed by integrating the full nonlinear SDEs in (7) under optimized control policies while taking into account the uncertainties introduced in Sec. IV.A.

V. Numerical Results

This section reports the numerical results of the nonlinear Monte Carlo simulations ($N_{\text{MC}} = 500$ samples) performed on the proposed chance-constrained trajectory optimization framework for the two representative transfers. The convex subproblems are modeled using YALMIP [24] and solved with the MOSEK solver [25]. The SCvx* parameters are the same as the ones reported in [21] for both covariance steering formulations investigated in this work.

A. Uncertainty Modeling

The operational uncertainties introduced in the nonlinear MC simulations include initial dispersion, OD error, maneuver execution error, and unmodeled dynamics modeled as a stochastic force. The initial state dispersion and initial estimate error are modeled as:

$$\hat{P}_0^- = \texttt{blkdiag}(\sigma_{r_0}^2 I_3, \sigma_{v_0}^2 I_3), \tag{23a}$$

$$\widetilde{P}_0^- = \texttt{blkdiag}(\sigma_{\text{obs},r}^2 I_3, \sigma_{\text{obs},v}^2 I_3) \tag{23b}$$

The target distribution in (11) and the maximum covariance in (12) follow an analogous structure.

The OD process is modeled as a full-state observation $\boldsymbol{f}_{\text{obs}}(\boldsymbol{x}_k) = \boldsymbol{x}_k$ with Gaussian noise: $G_{\text{obs}} = \texttt{blkdiag}(\sigma_{\text{obs},r} I_3, \sigma_{\text{obs},v} I_3)$ which represents the navigation or OD error. In the adopted setup, OD updates and impulsive maneuver opportunities are scheduled at the same discretization nodes, approximately once per day.

The stochastic acceleration is modeled as $G = [0_{3\times3}, \sigma_a I_3]^\top$ with $\sigma_a = 1 \times 10^{-11}$ km/s$^{3/2}$, consistent with previous CR3BP and planetary applications [16, 21]. In the MC validation, the nonlinear SDEs are integrated using a finer discretization, with each segment subdivided into 1000 intervals, yielding an effective white-noise update interval $\Delta t_{\text{WN}} \sim 1.5$ min in (8). For the maneuver execution error, only the proportional magnitude and pointing errors of the Gates model are considered. The uncertainty parameters are drawn from previous work on the Europa Clipper mission [11]; in this work, they are intentionally inflated with respect to the original values to highlight the approach robustness. Table 2 and Table 3 summarize the uncertainty and constraint parameters for the block Cholesky and full covariance formulations, respectively.

Table 2. Uncertainty and constraint parameters for BC

Quantity	Symbol	Value	Unit
Initial dispersion (pos.)	$\sigma_{r,0}$	100	m
Initial dispersion (vel.)	$\sigma_{v,0}$	10	mm/s
Initial est. error (pos.)	$\sigma_{\text{obs},r}$	3	m
Initial est. error (vel.)	$\sigma_{\text{obs},v}$	0.1	mm/s
Terminal dispersion (pos.)	$\sigma_{r,f}$	300	km
Terminal dispersion (vel.)	$\sigma_{v,f}$	10	m/s
Max. cov. (pos.)	$\sigma_{r,\text{max}}$	500	km
Max. cov. (vel.)	$\sigma_{v,\text{max}}$	100	m/s
Flyby max. cov. (pos.)	$\sigma_{r,\text{FB}}$	1-5	km
Flyby max. cov. (vel.)	$\sigma_{v,\text{FB}}$	0.1-1	m/s
Exec. error (prop. mag.)	σ_2	0.33	%
Exec. error (prop. point.)	σ_4	6.67	mrad
Stochastic acceleration	σ_a	10^{-11}	km/s$^{3/2}$
Max. control magnitude	u_{max}	10	m/s
Risk bound (control)	ϵ_u	10^{-3}	–
Risk bound (safety)	ϵ_{im}	10^{-3}	–

Table 3. Uncertainty and constraint parameters for FC

Quantity	Symbol	Value	Unit
Initial dispersion (pos.)	$\sigma_{r,0}$	100	m
Initial dispersion (vel.)	$\sigma_{v,0}$	10	mm/s
Initial est. error (pos.)	$\sigma_{\text{obs},r}$	30	m
Initial est. error (vel.)	$\sigma_{\text{obs},v}$	1	mm/s
Terminal dispersion (pos.)	$\sigma_{r,f}$	100	km
Terminal dispersion (vel.)	$\sigma_{v,f}$	10	m/s
Max. cov. (pos.)	$\sigma_{r,\text{max}}$	500	km
Max. cov. (vel.)	$\sigma_{v,\text{max}}$	100	m/s
Flyby max. cov. (pos.)	$\sigma_{r,\text{FB}}$	1	km
Flyby max. cov. (vel.)	$\sigma_{v,\text{FB}}$	0.1	m/s
Exec. error (prop. mag.)	σ_2	0.33	%
Exec. error (prop. point.)	σ_4	6.67	mrad
Stochastic acceleration	σ_a	10^{-11}	km/s$^{3/2}$
Max. control magnitude	u_{max}	5	m/s
Risk bound (control)	ϵ_u	10^{-3}	–
Risk bound (safety)	ϵ_{im}	10^{-3}	–

For the block Cholesky case, slightly looser bounds on flyby and terminal covariance constraints, and lower navigation uncertainty were required to achieve convergence.

B. Case 1: Non-resonant transfer

In the first scenario, block Cholesky and full covariance formulations are compared in terms of solution quality and constraints satisfaction. Fig. 3 and Fig. 5 report the MC trajectories (gray) with the target covariance (red) while Fig. 4 and Fig. 6 show the control maneuvers history throughout the transfer.

The comparison highlights a consistent advantage of the full covariance approach. It yields a lower ΔV budget and maintains tighter uncertainty control, whereas the block Cholesky formulation exhibits poorer covariance propagation, leading to larger dispersions both along the transfer and at the final node. The maximum covariance constraint is fundamental to keep the dispersion within a region where linearized uncertainty propagation remains meaningful. An alternative would be to implement the adaptive mechanism proposed in [26, 27] to regulate uncertainty growth by explicitly limiting local nonlinearity, which may improve the propagation of covariance, especially for the BC approach. Moreover, the full covariance formulation is significantly more computationally efficient, with a total solving time of about 1 minute compared to approximately 7 minutes for the block Cholesky approach.

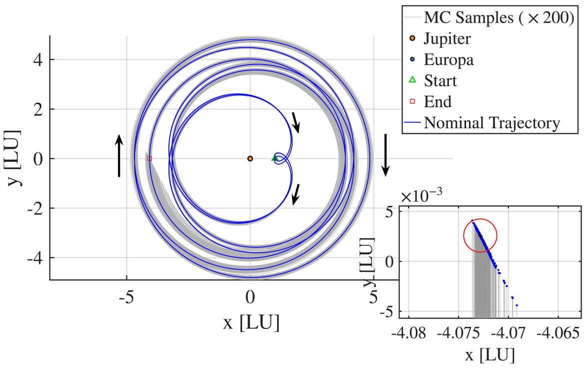

Fig. 3: *MC trajectories on xy-plane for block Cholesky*

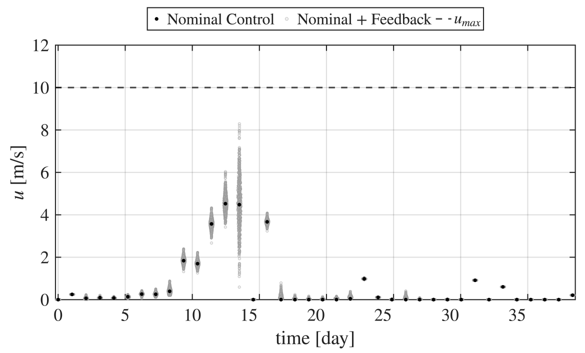

Fig. 4: *Control profile: nominal (black), MC samples (gray) for block Cholesky*

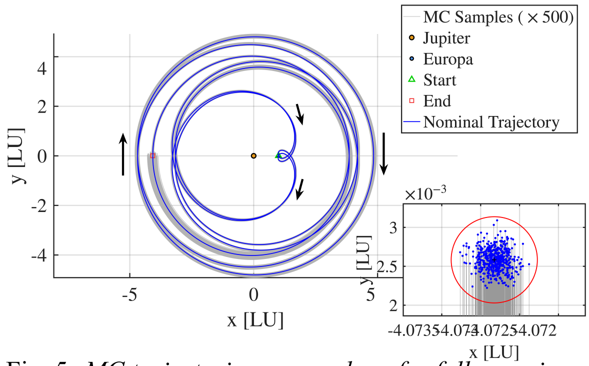

Fig. 5: *MC trajectories on xy-plane for full covariance*

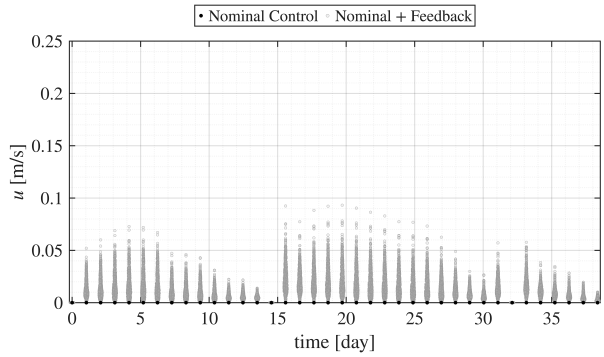

Fig. 6: *Control profile: nominal (black), MC samples (gray) for full covariance*

Finally, the results also reveal that both covariance steering formulations tend to produce dense maneuver schedules, which are undesirable. This motivates the hands-off covariance steering extension introduced in Section IV.E, whose results are reported in the following. Fig. 7 exhibits the sparser TCM schedule retrieved from the hands-off covariance steering optimization.

As shown in Fig. 8, the final node dispersion satisfies the prescribed target-covariance constraint: both with the predicted and actual 99.9% statistics, confirming that sparsity promotion does not compromise the assump-

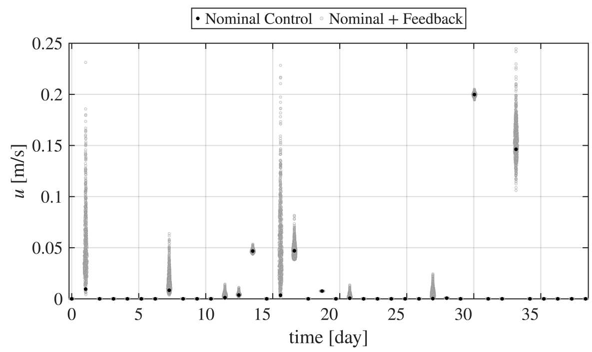

Fig. 7: *Control profile for hands-off covariance steering*

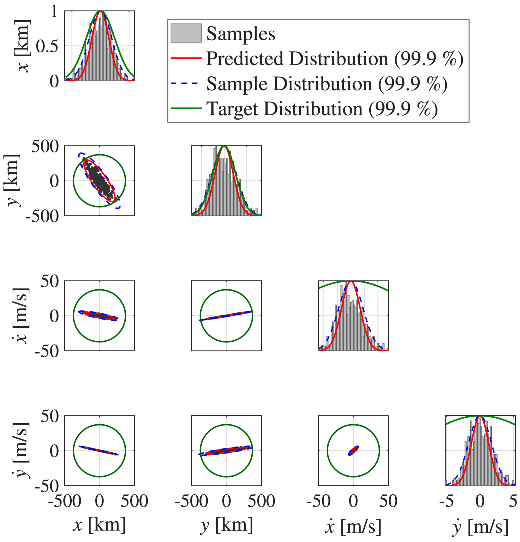

Fig. 8: *State dispersion at the final node for hands-off covariance steering*

tions on the linear covariance propagation.

The impact-avoidance role of the B-plane safety chance constraint is highlighted in Fig. 9 and 10. When the safety constraint is enforced in the optimization (Fig. 9), the predicted 99.9% covariance region and the Monte Carlo samples, propagated ballistically under nonlinear stochastic dynamics from the last maneuver before the assumed loss of control authority to the subsequent flyby, remain outside the impact boundary, confirming satisfaction of the prescribed PoI requirement. In contrast, when the safety constraint is removed in the optimization (Fig. 10), the distributions violate the impact surface, demonstrating that explicit inclusion of the B-plane safety constraint is essential to obtain a trajectory that is safe with respect to the Europa surface impact at flyby. The results are reported for the worst-case scenario of loss of control authority, corresponding to the failure occurring at the furthest node from the second Europa flyby, namely the first node where the safety constraint is applied.

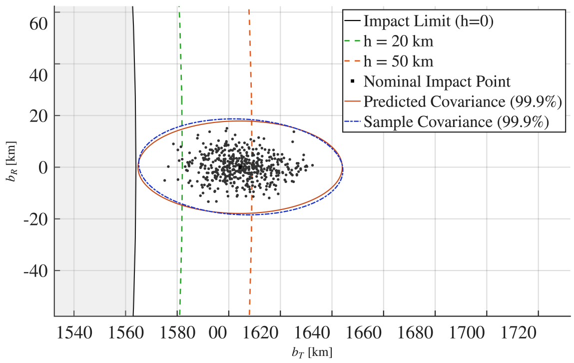

Fig. 9: *B-plane miss distance with the safety constraint with MC samples (black) at the second Europa flyby*

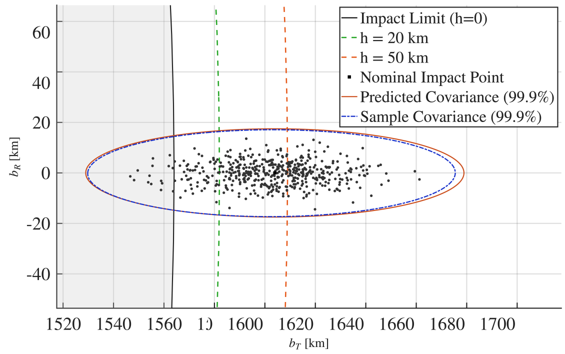

Fig. 10: *B-plane miss distance without the safety constraint at the second Europa flyby*

C. *Case 2: Resonant transfer*

For the resonant transfer case, the results are reported only for the full covariance formulation combined with sparsity promotion and are shown in Fig. 11, Fig. 12 and Fig. 13.

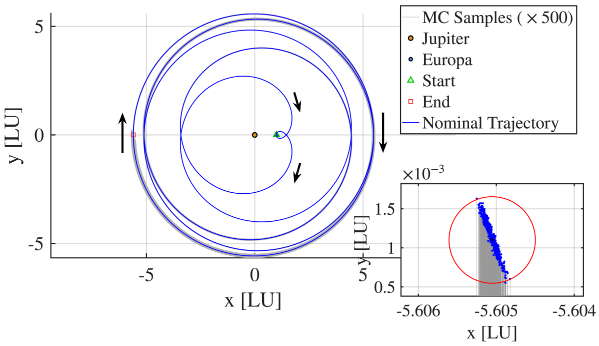

Fig. 11: *MC trajectories on xy-plane for the resonant transfer (full covariance with sparsity promotion)*

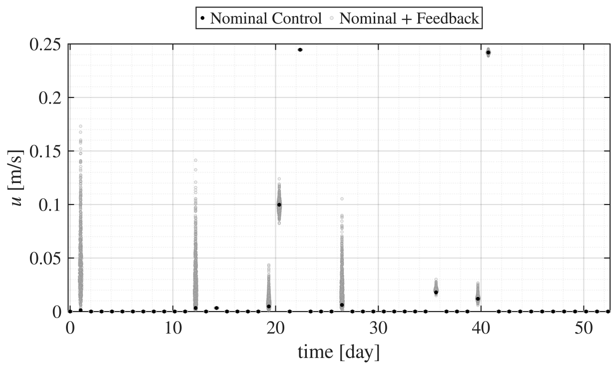

Fig. 12: *Control profile: nominal (black), MC samples (gray) for resonant transfer (full covariance with sparsity promotion)*

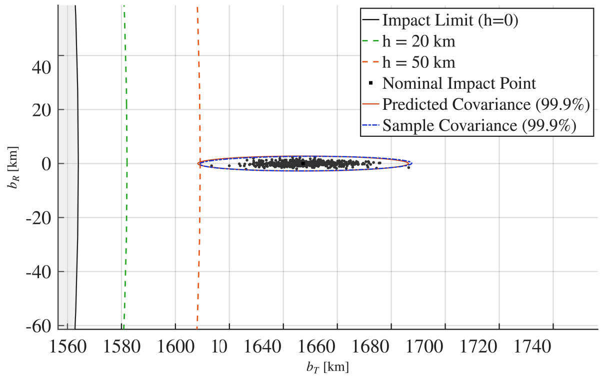

Fig. 13: *B-plane miss distance with the safety constraint with MC samples (black) at the second Europa flyby*

As in the previous case, the full covariance formulation with sparsity promotion successfully satisfies the terminal covariance and the B-plane safety constraints despite the longer time of flight and amplified sensitivity to uncertainties. The total ΔV_{99} is comparable to the one of the non-resonant case after sparsification. However, in the non-sparsified scenario, the resonant case presented a higher total ΔV_{99} mainly due to the increased feedforward effort required to enforce the safety constraint over the longer arcs. Sparsity promotion concentrates the correction effort into roughly four impulses per petal, yielding a greater sparsity level than in the non-resonant case while still reducing the total ΔV_{99} with respect to the non-sparsified solution.
Table 4 summarizes the main quantitative results across the two case studies, including the predicted and Monte Carlo ΔV_{99}, the number of active maneuvers, and the computational performance of each formulation. The $\Delta V_{99}^{\mathrm{UB}}$ is the convex upper bound on the 99th percentile of the closed-loop control effort, used as objective function in the convex subproblem 2; the MC ΔV_{99} is instead obtained empirically as the 99th percentile of the total ΔV across the Monte Carlo samples. As it can be noticed, the resulting ΔV_{99} of the BC approach is substantially larger than the one obtained with full covariance.

Table 4. Summary of the numerical results. BC stands for the block Cholesky approach, FC for the full covariance method, HO for hands-off covariance steering.

Case	Non res.	Non res.	Non res.	Res.
Formulation	BC	FC	FC + HO	FC + HO
$\Delta V_{99}^{\mathrm{UB}}$ [m/s]	39.13	1.75	1.13	1.21
$\Delta V_{99}^{\mathrm{MC}}$ [m/s]	33.40	1.51	0.88	0.94
Active TCMs	35	35	12	10
CPU time	$\sim$7 min	$\sim$1 min	$\sim$3 min	$\sim$3.5 min
SCvx* iter.	15	14	36	34

VI. Conclusions

This work demonstrates a chance-constrained trajectory design workflow for Europa multi-flyby transfers under realistic operational uncertainty. By combining sequential convex programming with covariance steering, the framework generated statistically robust nominal trajectories and feedback policies that satisfy B-plane targeting and PoI requirements, as confirmed through nonlinear Monte Carlo validation. While baseline solutions tended to have dense TCMs schedules, hands-off covariance steering effectively converted them into operationally plausible maneuver plans by concentrating the ΔV budget into fewer, well-separated TCMs while preserving probabilistic guarantees.
Overall, the proposed methodology provides a computationally tractable and practically effective approach for robust multi-flyby design under uncertainty. One of the most important direction for future work is to improve dispersion modeling in strongly nonlinear regimes moving beyond the Gaussian assumption by adopting non-Gaussian covariance steering approaches. Among others, possible developments are: to enhance the navigation process with a more realistic model, to move to higher-fidelity dynamical models, and to further refine the sparsity promotion algorithm. Further progress toward operational use will require improved verification and validation, including comparisons with higher-fidelity feedback strategies to better quantify performance and limitations.

VII. References

[1] M. G. Kivelson et al. Galileo magnetometer measurements: A stronger case for a subsurface ocean at europa. *Science*, 289(5483):1340–1343, 2000.

[2] L Alberto Cangahuala, Stefano Campagnola, Ben K Bradley, Dylan R Boone, Brent B Buffington, Jan M Ludwinski, Sumita Nandi, and Christopher J Scott. Europa clipper mission design, mission plan, and navigation. *Space Science Reviews*, 221(1):22, 2025.

[3] Jack Ridderhof, Kazuhide Okamoto, and Panagiotis Tsiotras. Chance constrained covariance control for linear stochastic systems with output feedback. In *2020 59th IEEE Conference on Decision and Control (CDC)*, pages 1758–1763. IEEE, 2020.

[4] Boris Benedikter, Alessandro Zavoli, Zhenbo Wang, Simone Pizzurro, and Enrico Cavallini. Convex approach to covariance control with application to stochastic low-thrust trajectory optimization. *Journal of Guidance, Control, and Dynamics*, 45(11):2061–2075, 2022.

[5] Kenshiro Oguri and Gregory Lantoine. Stochastic sequential convex programming for robust low-thrust trajectory design under uncertainty. In *AAS/AIAA Astrodynamics Specialist Conference*,

volume 8. AAS Charlotte, NC, 2022.

[6] Kenshiro Oguri. Chance-constrained control for safe spacecraft autonomy: Convex programming approach. In *2024 American Control Conference (ACC)*, pages 2318–2324. IEEE, 2024.

[7] Kazuhide Okamoto, Maxim Goldshtein, and Panagiotis Tsiotras. Optimal covariance control for stochastic systems under chance constraints. *IEEE Control Systems Letters*, 2(2):266–271, 2018.

[8] Naoya Kumagai and Kenshiro Oguri. Hands-off covariance steering: Inducing feedback sparsity via iteratively reweighted l1,p regularization. 2025. IEEE 64th Conference on Decision and Control (CDC).

[9] Aditya Arjun Anibha, Kenshiro Oguri, and Jacopo Villa. Autonomous optical navigation for unknown asteroids with irregular horizons. In *AAS/AIAA Astrodynamics Specialist Conference*, Boston, MA, 2025.

[10] Davide Farnocchia, Siegfried Eggl, Paul W Chodas, Jon D Giorgini, and Steven R Chesley. Planetary encounter analysis on the b-plane: a comprehensive formulation. *Celestial Mechanics and Dynamical Astronomy*, 131(8):36, 2019.

[11] Cristian Greco, Stefano Campagnola, and Massimiliano Vasile. Robust space trajectory design using belief optimal control. *Journal of Guidance, Control, and Dynamics*, 45(6):1060–1077, 2022.

[12] Thomas A Pavlak. *Trajectory design and orbit maintenance strategies in multi-body dynamical regimes*. PhD thesis, Purdue University, 2013.

[13] Stefano Campagnola, Brent B Buffington, Brian D Anderson, Etienne Pellegrini, Ricardo L Restrepo, Try Lam, Ben K Bradley, Christopher J Scott, Donald H Ellison, Kevin Bokelmann, et al. Europa clipper mission design: Design of the 21f31 reference tour. *The Journal of the Astronautical Sciences*, 72(6):64, 2025.

[14] Rodney L Anderson, Stefano Campagnola, and Brent B Buffington. Analysis of petal rotation trajectory characteristics. *Journal of Guidance, Control, and Dynamics*, 41(4):827–840, 2018.

[15] Kenshiro Oguri. *Risk-Aware Mission Design Around Small Celestial Bodies*. PhD thesis, University of Colorado at Boulder, 2021.

[16] Kenshiro Oguri and Gregory Lantoine. Convex block-cholesky approach to risk-constrained low-thrust trajectory design under operational uncertainty, 2026.

[17] Clarence R Gates. A simplified model of midcourse maneuver execution errors. Technical report, 1963.

[18] Kenshiro Oguri and Jay W McMahon. Stochastic primer vector for robust low-thrust trajectory design under uncertainty. *Journal of Guidance, Control, and Dynamics*, 45(1):84–102, 2022.

[19] Stefano Campagnola, Brent B Buffington, Try Lam, Mark Wallace, Troy Goodson, Andrew French, Dylan Boone, Zahi Tarzi, Narek R Shougarian, and Michael DiNicola. Europa clipper mission design: analysis of probability of impact requirements. *The Journal of the Astronautical Sciences*, 72(3):21, 2025.

[20] Naoya Kumagai and Kenshiro Oguri. Square root-factorized covariance steering. *arXiv preprint arXiv:2601.22348*, 2026.

[21] Naoya Kumagai and Kenshiro Oguri. Robust cislunar low-thrust trajectory optimization under uncertainties via sequential covariance steering. *Journal of Guidance, Control, and Dynamics*, 48:1–19, 08 2025.

[22] 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, 2023.

[23] Kenshiro Oguri and Gregory Lantoine. Lossless control-convex formulation for solar-sail trajectory optimization via sequential convex programming. *Journal of Guidance, Control, and Dynamics*, 48(2):311–326, 2025.

[24] J. Lofberg. Yalmip : a toolbox for modeling and optimization in matlab. In *2004 IEEE International Conference on Robotics and Automation (IEEE Cat. No.04CH37508)*, pages 284–289, 2004.

[25] MOSEK ApS. *The MOSEK Optimization Toolbox for MATLAB Manual, Version 11.3*. MOSEK ApS, 2026.

[26] Daniel C Qi and Kenshiro Oguri. Optimal linear covariance steering with minimum nonlinear dynamical errors. In *AAS/AIAA Astrodynamics Specialist Conference*, pages 1–20, 2025.

[27] L. Frassinella, K. Oguri, and F. Topputo. Safe autonomous guidance under uncertainty via sequential covariance steering for small-body proximity operations. In *30th International Symposium on Space Flight Dynamics (ISSFD)*, 2026.