

Revolution-Spaced Output-Feedback Model Predictive Control for Station Keeping on Near-Rectilinear Halo Orbits

Shimane, Yuri; Di Cairano, Stefano; Ho, Koki; Weiss, Avishai

TR2026-003 January 07, 2026

Abstract

We develop a model predictive control (MPC) policy for station keeping on a Near-Rectilinear Halo Orbit (NRHO). The proposed policy achieves full-state tracking of a reference NRHO via a multiple-maneuver control horizon, each spaced one revolution apart to abide by typical mission operation requirements. We prove that the proposed policy is recursively feasible, and perform numerical evaluation in an output-feedback setting by incorporating a navigation filter and realistic operational uncertainties, where the proposed MPC is compared against the state-of-the-art station-keeping algorithm adopted for the Gateway. Our approach successfully maintains the spacecraft in the vicinity of the reference NRHO at a similar cumulative cost as existing station keeping methods without encountering phase deviation issues, a common drawback of existing methods with one maneuver per revolution.

IEEE Transactions on Control Systems Technology 2025

© 2026 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Mitsubishi Electric Research Laboratories, Inc.
201 Broadway, Cambridge, Massachusetts 02139

Revolution-Spaced Output-Feedback Model Predictive Control for Station Keeping on Near-Rectilinear Halo Orbits

Yuri Shimane¹, Stefano Di Cairano², Koki Ho³, and Avishai Weiss⁴

Abstract—We develop a model predictive control (MPC) policy for station keeping on a Near-Rectilinear Halo Orbit (NRHO). The proposed policy achieves full-state tracking of a reference NRHO via a multiple-maneuver control horizon, each spaced one revolution apart to abide by typical mission operation requirements. We prove that the proposed policy is recursively feasible, and perform numerical evaluation in an output-feedback setting by incorporating a navigation filter and realistic operational uncertainties, where the proposed MPC is compared against the state-of-the-art station-keeping algorithm adopted for the Gateway. Our approach successfully maintains the spacecraft in the vicinity of the reference NRHO at a similar cumulative cost as existing station keeping methods without encountering phase deviation issues, a common drawback of existing methods with one maneuver per revolution.

I. INTRODUCTION

With growing interest in lunar exploration, *libration point orbits* (LPOs), quasi-periodic orbits about the equilibrium points of the Earth-Moon-spacecraft three-body system, offer unique locations to place both robotic and crewed spacecraft. For example, the lunar Gateway is planned in the 9:2 resonant southern Near-Rectilinear Halo Orbit (NRHO) about the Earth-Moon L2 point [1], [2], shown for one revolution in Figure 1. The instability of LPOs requires that the spacecraft conduct station-keeping (SK) maneuvers. The purpose of SK is to maintain the spacecraft near a pre-computed reference LPO, or *baseline*, in the presence of uncertainties such as state estimation error, modeling error, and control execution error. Due to the stringent propellant budget, typically higher instability of LPOs compared to traditional orbits around planets and moons, and the low number of heritage missions flying on LPOs, SK techniques on LPOs are an active area of research.

To accommodate mission operations, SK maneuvers are typically required to be as infrequent as possible [3]. On the NRHO with an orbital period of about 6.55 days, a typical requirement is for SK maneuvers to be conducted at most once every revolution about the Moon. To adhere to this requirement, a commonly adopted approach is known as *x-axis crossing control* (XAC) [3], a shooting-based method for designing SK maneuvers. Recently, the CAPSTONE mission [4] adopted XAC, and variants of XAC are currently being studied for the upcoming Gateway mission [3].

^{1,3}Y. Shimane and K. Ho are with the Daniel Guggenheim School of Aerospace Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA. Emails: {yuri.shimane, kokihoh} at gatech.edu

^{2,4}S. Di Cairano and A. Weiss are with Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139, USA. Emails: {dicairano, weiss} at merl.com

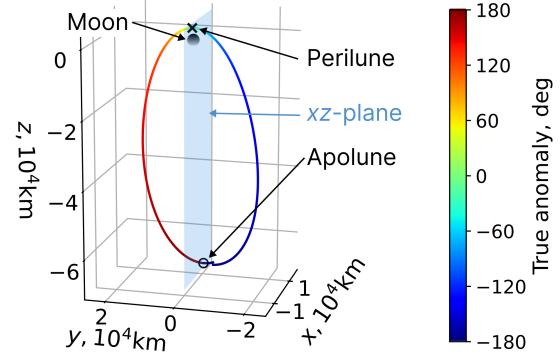


Fig. 1: NRHO in Earth-Moon rotating frame

One drawback of XAC stems from the fact that at most three out of the six translational state components can be assigned. To overcome this deficiency, XAC leverages the LPO's symmetry about the *xz*-plane in the Earth-Moon rotating frame [5], see Figure 1. A subset of the predicted spacecraft state at the intersection with the *xz*-plane is matched with the corresponding state components along the baseline when it intersects the same plane. Targeting based on *xz*-plane intersections results in a discrepancy between the epoch when the spacecraft crosses the plane and when the baseline crosses the plane. As a consequence, the steered path may experience a *phase angle deviation*: the spacecraft's location along the orbit may drift ahead or behind the baseline along the same orbit. To date, the phase deviation has been treated by ad-hoc heuristics, e.g., augmenting the targeting scheme with the epoch at which the symmetry event occurs [3], [6], or encapsulating the targeting scheme within a constrained optimization problem [7].

Beyond XAC, a number of model predictive control (MPC) policies have been proposed for SK [8], [9], [10], [11]. These policies do not suffer from phase angle deviation issues as they provide full-state tracking by considering either multiple maneuvers per revolution [9], [11] or a continuously controllable spacecraft [8], [10]. Moreover, the frequent control actions render these policies unfit for implementation in an actual mission. Some MPC approaches adopt quadratic cost functions [8], [11] that do not optimize the actual fuel, which is a scarce resource for a spacecraft. For further details, see [12] and references therein.

In this work, we propose an operationally compliant MPC policy that overcomes the phase disparity via full-state targeting. The proposed MPC uses a control horizon with at least two maneuvers spaced one revolution apart, which

provides sufficient controllability to track all six state components. Simultaneously, the one-revolution control cadence ensures that our approach is consistent with the operational requirement of conducting, at worst, a single SK maneuver per revolution. To explicitly minimize the propellant consumption, we employ an economic objective [13], [14] based solely on the control cost, resulting in similar SK cost to XAC over extended durations. The proposed MPC, hereafter denoted as SKMPC, sequentially solves a second-order cone program (SOCP) that steers the state of the spacecraft to the vicinity of the baseline at the end of its targeting horizon. At each iteration, the SOCP is re-instantiated by linearizing the dynamics about the steered state from the previous iteration; the SKMPC is terminated when the steered state propagated with the nonlinear dynamics lies sufficiently close to the baseline. We provide a brief discussion on the recursive feasibility of the SKMPC and numerically demonstrate its performance, with comparison to XAC.

This work extends [15] by incorporating a navigation filter to estimate the full state of the spacecraft, validating the proposed approach in a realistic output-feedback scenario. Our simulation incorporates disturbances due to navigational uncertainty, dynamics modeling errors, control actuation errors, and random impulses imparted at scheduled times along the NRHO due to momentum wheel desaturation maneuvers. We also implement XAC and compare its performance against the SKMPC. We provide comprehensive Monte Carlo results with varying disturbance levels, thereby quantifying the coupled performance of the filter and the controller.

II. BACKGROUND

A. Spacecraft Dynamics Model

We consider the spacecraft's motion in the J2000 inertial frame \mathcal{F}_{Inr} [16], centered at the Moon. The state of the spacecraft $\mathbf{x} \in \mathbb{R}^6$ consists of the Cartesian position $\mathbf{r} \in \mathbb{R}^3$ with respect to the Moon and the rate of change of \mathbf{r} in \mathcal{F}_{Inr} , denoted by $\mathbf{v} \in \mathbb{R}^3$. The equations of motion are [17]

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t) = \begin{bmatrix} \mathbf{v} \\ -\frac{\mu}{r^3} \mathbf{r} + \mathbf{a}_{\text{J2}} + \sum_i \mathbf{a}_{N_i} + \mathbf{a}_{\text{SRP}} \end{bmatrix}, \quad (1)$$

where $r = \|\mathbf{r}\|_2$, and μ is the gravitational parameter of the Moon. The derivative of \mathbf{v} consists, in order, of the Keplerian acceleration due to the Moon, J2 perturbation of the Moon \mathbf{a}_{J2} , gravitational perturbations by other celestial bodies \mathbf{a}_{N_i} , and the solar radiation pressure (SRP) \mathbf{a}_{SRP} , all in \mathcal{F}_{Inr} . We include third-body perturbations due to the Earth and the Sun. For expressions of \mathbf{a}_{J2} , \mathbf{a}_{N_i} , and \mathbf{a}_{SRP} , see [17]. Note that \mathbf{a}_{J2} , \mathbf{a}_{N_i} and \mathbf{a}_{SRP} in equation (1) are time-dependent, making \mathbf{f} non-autonomous. Constants in the equations of motion and ephemerides of celestial bodies are obtained from the SPICE toolkit [16].

An initial perturbation $\delta\mathbf{x}_k$ at time t_k can be linearly propagated to time t_{k+1} , denoted by $\delta\mathbf{x}_{k+1}$, via the linear state-transition matrix (STM) $\Phi(t_{k+1}, t_k) \in \mathbb{R}^{6 \times 6}$,

$$\delta\mathbf{x}_{k+1} = \Phi(t_{k+1}, t_k) \delta\mathbf{x}_k. \quad (2)$$

The Jacobian of the dynamics may be used to construct the STM by solving the matrix initial value problem (IVP)

$$\dot{\Phi}(t, t_k) = \frac{\partial \mathbf{f}(\mathbf{x}, t)}{\partial \mathbf{x}} \Phi(t, t_k), \quad \Phi(t_k, t_k) = \mathbf{I}_6. \quad (3)$$

We use the shorthand notations $\mathbf{x}_k = \mathbf{x}(t_k)$ and $\Phi_{j,i} = \Phi(t_j, t_i)$, and we express the block submatrices of $\Phi_{j,i}$ as

$$\Phi_{j,i} = \begin{bmatrix} \Phi_{j,i}^{rr} & \Phi_{j,i}^{rv} \\ \Phi_{j,i}^{vr} & \Phi_{j,i}^{vv} \end{bmatrix}.$$

Assuming the control on the spacecraft results in an impulsive change in velocity¹, the state at time t_{k+1} following an impulse applied at time t_k is given by

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \int_{t_k}^{t_{k+1}} \mathbf{f}(\mathbf{x}, t) + \delta(t - t_k) \begin{bmatrix} \mathbf{0}_{3 \times 1} \\ \mathbf{u}_k \end{bmatrix} dt, \quad (4)$$

where the control $\mathbf{u}_k \in \mathbb{R}^3$ is the impulsive change in velocity and δ is the Dirac delta function.

B. Earth-Moon Rotating Frame

The Earth-Moon rotating frame, \mathcal{F}_{EM} , is practical for analyzing the natural motion and devise SK controllers due to the NRHO's symmetry about the frame's xz -plane, as illustrated in Figure 1. In \mathcal{F}_{EM} , the x -axis is aligned with the instantaneous Earth-Moon vector, the z -axis is aligned with the co-rotating angular velocity vector of the Earth and the Moon, and the y -axis completes the triad. The matrix $\mathbf{S}_{\text{EM}}^{\text{Inr}} \in \mathbb{R}^{6 \times 6}$ that transforms a state from \mathcal{F}_{Inr} to \mathcal{F}_{EM} is

$$\mathbf{S}_{\text{EM}}^{\text{Inr}}(t) = \begin{bmatrix} \mathbf{T}_{\text{EM}}^{\text{Inr}}(t) & \mathbf{0}_{3 \times 3} \\ \dot{\mathbf{T}}_{\text{EM}}^{\text{Inr}}(t) & \mathbf{T}_{\text{EM}}^{\text{Inr}}(t) \end{bmatrix}, \quad \mathbf{T}_{\text{EM}}^{\text{Inr}} = \begin{bmatrix} \mathbf{e}_1^T(t) \\ \mathbf{e}_2^T(t) \\ \mathbf{e}_3^T(t) \end{bmatrix}, \quad (5)$$

where $\mathbf{e}_1(t)$, $\mathbf{e}_2(t)$, and $\mathbf{e}_3(t)$ are the basis vectors of \mathcal{F}_{EM} realized in \mathcal{F}_{Inr} ,

$$\mathbf{e}_1(t) = \frac{-\mathbf{d}_{\text{Earth}}(t)}{\|\mathbf{d}_{\text{Earth}}(t)\|_2}, \quad \mathbf{e}_3(t) = \frac{\mathbf{d}_{\text{Earth}}(t) \times \mathbf{v}_{\text{Earth}}(t)}{\|\mathbf{d}_{\text{Earth}}(t) \times \mathbf{v}_{\text{Earth}}(t)\|_2},$$

and $\mathbf{e}_2(t) = \mathbf{e}_3(t) \times \mathbf{e}_1(t)$, $\mathbf{v}_{\text{Earth}}(t)$ is the Earth's velocity with respect to the Moon, and $\dot{\mathbf{T}}_{\text{EM}}^{\text{Inr}}(t)$ is the time-derivative of $\mathbf{T}_{\text{EM}}^{\text{Inr}}(t)$ in \mathcal{F}_{Inr} .

C. Stability on Near Rectilinear Halo Orbit

The NRHO's unstable subspace [2] necessitates SK actions to prevent the spacecraft from diverging from the baseline in the presence of uncertainties. Let θ denote the *osculating true anomaly* with respect to the Moon [17],

$$\theta = \text{atan2}(\mathbf{h} \mathbf{v}_r, h^2/r - \mu), \quad h = \|\mathbf{r} \times \mathbf{v}\|_2, \quad \mathbf{v}_r = \frac{\mathbf{r} \cdot \mathbf{v}}{r}.$$

As noted in the literature on NRHO [2], the dynamics are most sensitive around *perilune* where $\theta = 0^\circ$ and the spacecraft is closest to the Moon, and least sensitive around *apolune* where $\theta = 180^\circ$ and the spacecraft is farthest from the Moon. To make the SK activity as robust as possible against navigation and control actuation errors, SK

¹Due to control executions lasting on the order of seconds to minutes along an orbit with a period on the order of days, all conventional thrusters are effectively impulsive in this application [17].

maneuvers typically execute near apolune. Let the *maneuver true anomaly* θ_{man} denote the true anomaly at which the SK maneuver is scheduled to occur. In accordance with operational plans for the Gateway [3], we use $\theta_{\text{man}} = 200^\circ$.

D. Navigation Filter

We consider an extended Kalman filter (EKF) to estimate the spacecraft state. Let $\hat{\mathbf{x}} \in \mathbb{R}^6$ and $\mathbf{P} \in \mathbb{R}^{6 \times 6}$ denote the estimate of the state \mathbf{x} and the state covariance, respectively. We briefly present the prediction and update steps of the EKF, along with the measurement model and the impulse events, where predicted quantities at t_k based on t_{k-1} are denoted by $(\cdot)_k^-$, and updated quantities following a measurement or a maneuver at t_k are denoted by $(\cdot)_k^+$.

1) *Prediction*: The prediction step from t_{k-1} to t_k is [18]

$$\begin{aligned}\hat{\mathbf{x}}_k^- &= \hat{\mathbf{x}}_{k-1}^+ + \int_{t_{k-1}}^{t_k} \mathbf{f}(\hat{\mathbf{x}}, t) dt, \\ \mathbf{P}_k^- &= \Phi_{k,k-1} \mathbf{P}_{k-1}^+ \Phi_{k,k-1}^T + \mathbf{Q}_{k,k-1},\end{aligned}$$

where $\Phi_{k,k-1}$ is the STM evaluated along $\hat{\mathbf{x}}_k^-$, and $\mathbf{Q}_{k,k-1}$ is the process noise accounting for unmodelled disturbances. We adopt the unbiased random process noise model [19]

$$\mathbf{Q}_{k,k-1} = \sigma_p^2 \begin{bmatrix} (\Delta t^3/3) \mathbf{I}_3 & (\Delta t^2/2) \mathbf{I}_3 \\ (\Delta t^2/2) \mathbf{I}_3 & \Delta t \mathbf{I}_3 \end{bmatrix}, \quad \Delta t = t_k - t_{k-1}$$

where σ_p is a tuning parameter.

2) *Update*: At time t_k , a measurement is provided to the filter. The noisy measurement $\mathbf{y}_k \in \mathbb{R}^m$ is assumed to follow a multivariate normal distribution with zero mean and covariance \mathbf{R}_k [18] such that $\mathbf{y}_k = \mathbf{h}(\hat{\mathbf{x}}_k^-) + \mathcal{N}(\mathbf{0}_{m \times 1}, \mathbf{R}_k)$. Let $\mathbf{H}_k = \mathbf{H}(\mathbf{x}_k) = \partial \mathbf{h}(\hat{\mathbf{x}}_k^-) / \partial \mathbf{x}$; the update step is [18]

$$\begin{aligned}\mathbf{L}_k &= \mathbf{P}_k^- \mathbf{H}_k^T \left(\mathbf{H}_k \mathbf{P}_k^- \mathbf{H}_k^T + \mathbf{R}_k \right)^{-1}, \\ \hat{\mathbf{x}}_k^+ &= \hat{\mathbf{x}}_k^- + \mathbf{L}_k (\mathbf{y}_k - \mathbf{h}_k(\hat{\mathbf{x}}_k^-)), \\ \mathbf{P}_k^+ &= (\mathbf{I}_6 - \mathbf{L}_k \mathbf{H}_k) \mathbf{P}_k^- (\mathbf{I}_6 - \mathbf{L}_k \mathbf{H}_k)^T + \mathbf{L}_k \mathbf{R}_k \mathbf{L}_k^T.\end{aligned}$$

The Joseph form of the covariance update has been adopted for superior numerical stability [19].

3) *Measurements*: We consider measurements based on range and range-rate, with \mathbf{h} and \mathbf{H} given by

$$\begin{aligned}\mathbf{h}(\mathbf{x}) &= \begin{bmatrix} r \\ \dot{r} \end{bmatrix} = \begin{bmatrix} \|\mathbf{r}\|_2 \\ \mathbf{r}^T \mathbf{v} / \|\mathbf{r}\|_2 \end{bmatrix}, \\ \mathbf{H}(\mathbf{x}) &= \begin{bmatrix} \frac{x}{r} & \frac{y}{r} & \frac{z}{r} & 0 & 0 & 0 \\ \frac{v_x}{r} - \frac{x\dot{r}}{r^2} & \frac{v_y}{r} - \frac{y\dot{r}}{r^2} & \frac{v_z}{r} - \frac{z\dot{r}}{r^2} & \frac{x}{r} & \frac{y}{r} & \frac{z}{r} \end{bmatrix},\end{aligned}$$

where \mathbf{r} , \mathbf{v} , and the state components $[x, y, z, v_x, v_y, v_z]$ are in \mathcal{F}_{Inr} . We assume a constant measurement covariance $\mathbf{R}_k = \mathbf{R} = \text{diag}(\sigma_r^2, \sigma_{\dot{r}}^2)$ where σ_r and $\sigma_{\dot{r}}$ are the standard deviations of the range and range-rate measurements.

4) *Impulse Events*: We model SK maneuvers as a velocity impulse on the spacecraft $\mathbf{u}_k = \hat{\mathbf{u}}_k + \mathcal{N}(\mathbf{0}_{3 \times 1}, \mathbf{V}_k)$, where $\hat{\mathbf{u}}$ is the expected impulse, and \mathbf{V}_k is the corresponding covariance. The maneuver estimate $\hat{\mathbf{u}}$ is computed by the SK controller and $\mathbf{V}_k = (\sigma_{\mathbf{u},\text{abs}} + \sigma_{\mathbf{u},\text{rel}} \|\hat{\mathbf{u}}\|_2)^2 \mathbf{I}_3$, where

$\sigma_{\mathbf{u},\text{abs}}$ and $\sigma_{\mathbf{u},\text{rel}}$ are the absolute and relative standard deviation of the thruster. Then, the state is updated via $\hat{\mathbf{x}}_k^+ = \hat{\mathbf{x}}_k^- + [\mathbf{0}_{3 \times 1}^T, \hat{\mathbf{u}}^T]^T$, and the covariance is updated via

$$\mathbf{P}_k^+ = \mathbf{P}_k^- + \begin{bmatrix} \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{V}_k \end{bmatrix}.$$

Such an update following an uncertain maneuver effectively works as a partial reset to the EKF. Had we not assumed an impulsive control model, we would need to consider a compensation model with intensities scaled according to the maneuver magnitude [19].

E. x-Axis Crossing Control

We now provide a brief review of XAC. XAC attempts to maintain the symmetry of the controlled trajectory about the xz -plane in \mathcal{F}_{EM} [3]. At a given control time t_k , let t_{Np} and $t_{Np,\text{ref}}$ denote the $N_{\text{rev}}^{\text{th}}$ times the predicted trajectory and the baseline cross the xz -plane near perilune, respectively, and $\mathbf{x}_{Np,\text{ref}}$ denote the baseline state at $t_{Np,\text{ref}}$. XAC is a single shooting problem to obtain a control \mathbf{u}_k that results in residual F defined as the difference between the actual and the baseline x -axis component of the velocity in \mathcal{F}_{EM} when crossing the xz -plane,

$$F = \mathbf{\Lambda} \left(\mathbf{S}_{\text{EM}}^{\text{Inr}}(t_{Np}) \mathbf{x}_{Np} - \mathbf{S}_{\text{EM}}^{\text{Inr}}(t_{Np,\text{ref}}) \mathbf{x}_{Np,\text{ref}} \right),$$

$$\mathbf{\Lambda} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix},$$

$$\mathbf{x}_{Np} = \hat{\mathbf{x}}_k^- + \begin{bmatrix} \mathbf{0}_{3 \times 1} \\ \mathbf{u}_k \end{bmatrix} + \int_{t_k}^{t_{Np}} \mathbf{f}[\mathbf{x}(t), t] dt,$$

being below a user-defined threshold ϵ_F . Note that $t_{Np} \neq t_{Np,\text{ref}}$ since these are the times at the xz -plane crossing along two different trajectories. The maneuver $\mathbf{u}(t_k)$ is obtained by Newton-Raphson iterations,

$$\mathbf{u}_k \leftarrow \mathbf{u}_k - \mathbf{D}F^T (\mathbf{D}F \mathbf{D}F^T)^{-1} F(\mathbf{u}_k), \quad \mathbf{D}F = \frac{\partial F}{\partial \mathbf{u}_k}, \quad (10)$$

until $|F| \leq \epsilon_{v_x}$. Since XAC suffers from phase dispersion, Davis et al. [3] propose augmenting F with the difference between t_{Np} and $t_{Np,\text{ref}}$; however, this approach requires balancing physically different quantities of F during the Newton-update, making it sensitive to hyperparameters and non-intuitive to tune [7].

III. FULL-STATE TARGETING ECONOMIC MPC FOR STATION-KEEPING ON NRHO

SKMPC computes a sequence of maneuvers based on the state estimate $\hat{\mathbf{x}}_k^-$ at the current time t_k to ensure the predicted state at some future time $t_f > t_k$ lies in the vicinity of the baseline. To reduce the sensitivity of the errors on the targeted state, we select t_f to correspond to the apolune along the baseline approximately N_{rev} revolutions in the future. This proposed method computes an SK maneuver at the prescribed θ_{man} such that the steered state lies in the vicinity of the baseline at the $N_{\text{rev}}^{\text{th}}$ apolune into the future.

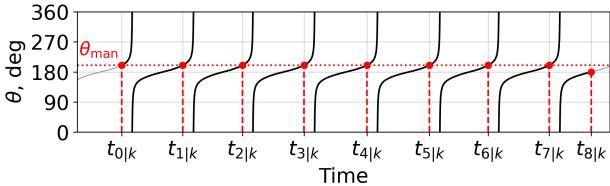


Fig. 2: Control horizon with $N = 9$ and $\theta_{\text{man}} = 200^\circ$

A. Problem Formulation

Let $(\cdot)_{j|k}$ denote a quantity predicted at j time increments ahead of the EKF's estimate $\hat{\mathbf{x}}_k^-$ at t_k . Let $N = N_{\text{rev}} + 1$ denote the number of impulse maneuvers in the control horizon, with controls denoted by $\mathbf{u}_{j|k} \in \mathbb{R}^3$, $j = 0, \dots, N-1$, occurring at times

$$t_{j|k} = \begin{cases} t(\theta_{\text{man}}, j) & j \leq N-2, \\ t(180^\circ, j) & j = N-1, \end{cases} \quad (11)$$

where $t(a, b)$ denote the earliest time $t \geq t_k$ such that $\theta(t) = a$ after completing b revolutions. Figure 2 illustrates the control horizon defined by (11) for $N = 9$ and $\theta_{\text{man}} = 200^\circ$, corresponding to about 50 days.

We formulate a minimization problem with an economic sum-of-2-norm objective of N maneuvers, which corresponds directly to the propellant mass consumed via Tsolkovsky's rocket equation [17]. Let $\mathbf{X} \in \mathbb{R}^{6 \times N}$ and $\mathbf{U} \in \mathbb{R}^{3 \times N}$ denote the concatenated states and controls

$$\mathbf{X} = [\mathbf{x}_{0|k} \ \cdots \ \mathbf{x}_{N-1|k}], \quad \mathbf{U} = [\mathbf{u}_{0|k} \ \cdots \ \mathbf{u}_{N-1|k}].$$

Let $\bar{\mathbf{x}}_{j|k}$ denote the predicted trajectory prior to computing the SK maneuvers, and $\bar{\mathbf{u}}_{j|k}$ denote pre-planned maneuvers, initialized to $\mathbf{0}_{3 \times 1}$. We define $\bar{\mathbf{X}}$ and $\bar{\mathbf{U}}$ as

$$\bar{\mathbf{X}} = [\bar{\mathbf{x}}_{0|k} \ \cdots \ \bar{\mathbf{x}}_{N-1|k}], \quad \bar{\mathbf{U}} = [\bar{\mathbf{u}}_{0|k} \ \cdots \ \bar{\mathbf{u}}_{N-1|k}].$$

We begin by defining the discrete-time optimal control problem (OCP) assuming linearized dynamics around $\bar{\mathbf{x}}$, then will subsequently consider a sequential linearization scheme in Section III-D, where the OCP will be solved repeatedly, each time updating $\bar{\mathbf{X}}$ and $\bar{\mathbf{U}}$. The OCP with linearized dynamics is

$$\min_{\mathbf{X}, \mathbf{U}} \sum_{j=0}^{N-1} \|\mathbf{u}_{j|k}\|_2 \quad (12a)$$

$$\text{s.t. } \mathbf{x}_{j+1|k} = \Phi_{j+1,j|k} \mathbf{x}_{j|k} + \begin{bmatrix} \Phi_{j+1,j|k}^{rv} \\ \Phi_{j+1,j|k}^{vv} \end{bmatrix} \mathbf{u}_{j|k} + \mathbf{c}_{j+1,j|k} \quad (12b)$$

$$j = 0, \dots, N-2$$

$$\mathbf{x}_{0|k} = \hat{\mathbf{x}}_k^- \quad (12c)$$

$$\mathbf{x}_{N-1|k} + \begin{bmatrix} \mathbf{0}_{3 \times 1} \\ \mathbf{u}_{N-1|k} \end{bmatrix} \in \mathcal{X}_{\text{targ}}(t_{N-1|k}) \quad (12d)$$

$$\mathbf{x}_{j|k} \in \mathcal{X}, \quad \forall j = 0, \dots, N-1 \quad (12e)$$

$$\mathbf{u}_{j|k} \in \mathcal{U}, \quad \forall j = 0, \dots, N-1, \quad (12f)$$

where \mathcal{X} is the admissible set for the state, \mathcal{U} is the admissible set for the control, $\Phi_{j+1,j|k}$ is the STM from $t_{j|k}$

to $t_{j+1|k}$ with initial condition $\bar{\mathbf{x}}_{j|k}$ and impulsive control $\bar{\mathbf{u}}_{j|k}$, and $\mathbf{c}_{j+1,j|k}$ is given by

$$\mathbf{c}_{j+1,j|k} = \bar{\mathbf{x}}_{j+1,j|k} - \Phi_{j+1,j|k} \left(\bar{\mathbf{x}}_{j|k} + \begin{bmatrix} \mathbf{0}_{3 \times 1} \\ \bar{\mathbf{u}}_{j|k} \end{bmatrix} \right), \quad (13)$$

$$\bar{\mathbf{x}}_{j+1,j|k} = \bar{\mathbf{x}}_{j|k} + \int_{t_{j|k}}^{t_{j+1|k}} \mathbf{f}(\bar{\mathbf{x}}, t) + \delta(t - t_k) \begin{bmatrix} \mathbf{0}_{3 \times 1} \\ \bar{\mathbf{u}}_{j|k} \end{bmatrix} dt.$$

The linearized dynamics in (12b) implies that the control action $\mathbf{u}_{j|k}$ shifts the state within some trust-region $\delta \in \mathbb{R}^6$,

$$|\mathbf{x}_{j|k} - \bar{\mathbf{x}}_{j|k}| \leq \delta, \quad j = 0, \dots, N-1.$$

Constraint (12c) enforces the initial state to coincide with the current state estimate, and constraint (12d) enforces the final state to lie in $\mathcal{X}_{\text{targ}}(t_{N-1|k})$. In (12f), \mathcal{U} is defined with a maximum executable control magnitude u_{max} as

$$\mathcal{U} = \{\mathbf{u} \in \mathbb{R}^3 : \|\mathbf{u}\|_2 \leq u_{\text{max}}\}. \quad (14)$$

B. Definition of Terminal Constraint Set

Let $\mathcal{X}_{\text{targ}}(t_{N-1|k})$ be a 6D ellipsoid centered at the baseline state $\mathbf{x}_{f,\text{ref}} \triangleq [\mathbf{r}_{f,\text{ref}}^T, \mathbf{v}_{f,\text{ref}}^T]^T$ in \mathcal{F}_{EM} at time $t_{N-1|k}$,

$$\mathcal{X}_{\text{targ}}(t_{N-1|k}) = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{r} - \mathbf{r}_{f,\text{ref}}\|_2 \leq \epsilon_r, \|\mathbf{v} - \mathbf{v}_{f,\text{ref}}\|_2 \leq \epsilon_v\}, \quad (15)$$

where ϵ_r and ϵ_v are the magnitudes of the apses of the ellipsoid along position and velocity components, and are tuning parameters. The terminal constraint (12d) can be replaced by two second-order cone (SOC) constraints,

$$\left\| \mathbf{r}_{N-1|k}^{\text{EM}} - \mathbf{r}_{f,\text{ref}} \right\|_2 \leq \epsilon_r, \quad (16a)$$

$$\left\| \mathbf{v}_{N-1|k}^{\text{EM}} + \mathbf{T}_{\text{EM}}^{\text{Inr}} \mathbf{u}_{N-1|k} - \mathbf{v}_{f,\text{ref}} \right\|_2 \leq \epsilon_v, \quad (16b)$$

where $\mathbf{r}_{N-1|k}^{\text{EM}} = \mathbf{T}_{\text{EM}}^{\text{Inr}}(t_{N-1|k}) \mathbf{r}_{N-1|k}$ and $\mathbf{v}_{N-1|k}^{\text{EM}} = \mathbf{T}_{\text{EM}}^{\text{Inr}}(t_{N-1|k}) \mathbf{v}_{N-1|k} + \mathbf{T}_{\text{EM}}^{\text{Inr}}(t_{N-1|k}) \mathbf{u}_{N-1|k}$.

C. Recursive Feasibility

Next, we briefly discuss the recursive feasibility of problem (12) with input constraint (14) and terminal set constraint (15). The non-autonomous dynamics (1) results in the terminal set $\mathcal{X}_{\text{targ}}(t_{N-1|k})$ to also be time-dependent, complicating the application of the classical approach for proving recursive feasibility of MPC. Computing and storing such a time-varying set for the NRHO, which is not periodic but only quasi-periodic, may be prohibitive in practice. However, this specific application has some favorable conditions that help us recover guarantees of recursive feasibility. First, for the considered family of orbits, the STM in (12b) ensures controllability of the linearized system around the nominal orbit, described by $\mathbf{f}(\mathbf{x}_k, t)$. Second, the available thrust upper-bounded by u_{max} is “significantly larger” than what is required in SK maneuvers, although the general desire is to minimize the requested thrust.

Proposition 1: Let $2 \leq K \leq N$ correspond to the number of maneuvers such that

$$\text{rank} \left(\left[\begin{bmatrix} \Phi_{N-1,0|k}^{rv} \\ \Phi_{N-1,0|k}^{vv} \end{bmatrix} \dots \begin{bmatrix} \Phi_{N-1,K-1|k}^{rv} \\ \Phi_{N-1,K-1|k}^{vv} \end{bmatrix} \right] \right) = 6.$$

For a large enough u_{\max} , if (12) is feasible at time t_{k-1} then it is feasible at t_k . Furthermore, the trajectory remains bounded in a set \mathcal{X}_{bnd} at the apolune times t_k .

Proof. Since (12) is feasible at t_{k-1} , there exists $U_K(t_{k-1}) = [\mathbf{u}_{0|k-1}, \dots, \mathbf{u}_{K-1|k-1}]$ such that $\mathbf{x}_{N-1|k-1} \in \mathcal{X}_{\text{targ}}(t_{N-1|k-1})$, where $\mathbf{x}_{N-1|k-1}$ is obtained by applying the entire sequence $U_K(t_{k-1})$ via (12b) followed by open loop evolution from $t_{K-1|k-1}$ to $t_{N-1|k-1}$. Let $\mathbf{x}_{N|k-1} = \mathbf{x}_{N-1|k-1} + \int_{t_{N-1|k-1}}^{t_{N-1|k}} \mathbf{f}(\mathbf{x}, t) dt$. We need to prove that it is possible to obtain a state perturbation $\delta\mathbf{x}_{N-1|k}$ such that $\mathbf{x}_{N|k-1} + \delta\mathbf{x}_{N-1|k} \in \mathcal{X}_{\text{targ}}(t_{N-1|k})$.

Consider the candidate control sequence $U_K(t_k) = [\mathbf{u}_{1|k-1} + \delta\mathbf{u}_{0|k}, \dots, \mathbf{u}_{K-1|k-1} + \delta\mathbf{u}_{K-2|k}, \mathbf{u}_{K-1|k}]$ and

$$\begin{aligned} & \delta\mathbf{x}_{N-1|k} \\ &= \sum_{j=0}^{K-2} \left[\begin{bmatrix} \Phi_{N-1,j|k}^{\mathbf{rv}} \\ \Phi_{N-1,j|k}^{\mathbf{vv}} \end{bmatrix} \delta\mathbf{u}_{j|k} + \begin{bmatrix} \Phi_{N-1,K-1|k}^{\mathbf{rv}} \\ \Phi_{N-1,K-1|k}^{\mathbf{vv}} \end{bmatrix} \mathbf{u}_{K-1|k} \right]. \end{aligned}$$

Then, $\mathbf{x}_{N-1|k} = \mathbf{x}_{N|k-1} + \delta\mathbf{x}_{N-1|k} \in \mathcal{X}_{\text{targ}}(t_{N-1|k})$ is guaranteed by the controllability in K steps for some $\Delta U_K(t_k) = [\delta\mathbf{u}_{0|k}, \dots, \delta\mathbf{u}_{K-2|k}, \mathbf{u}_{K-1|k}]$ which perturbs and extends the previous control sequence $U_K(t_{k-1})$. For a large enough u_{\max} , the sequence $U_K(t_k)$ is feasible. Due to the finite horizon and the bounded thrust, the trajectories remain bounded in a set \mathcal{X}_{bnd} because the control strategy enforces (12d), and $\mathcal{X}_{\text{targ}}(t_{N-1|k})$ is bounded by (15). ■

With regards to the assumptions in Proposition 1, due to the quasi-periodic nature of the orbit, the difference between $\mathcal{X}_{\text{targ}}(t_{N-1|k-1})$ and $\mathcal{X}_{\text{targ}}(t_{N-1|k})$ is usually small. The necessitated correction $\delta\mathbf{x}$ is thus relatively small compared to the control authority u_{\max} . Hence, the maximum thrust of the propulsion system will be sufficient to ensure the feasibility of the candidate control sequence. The controllability of the spacecraft in the NRHO orbit ensures the rank condition.

The economic structure of the SKMPC, along with its operational setting, complicates the theoretical demonstration of closed-loop stability. Angeli [14] demonstrates closed-loop stability for economic MPC operating with respect to an equilibrium point for a strictly dissipative system by introducing a dissipation inequality. With the present SKMPC, the control is with respect to a quasi-periodic orbit in non-autonomous dynamics and in the presence of uncertainties. Guaranteeing closed-loop stability would require modifying the OCP (12), which would compromise its fuel-optimal structure [20], and it is not clear whether stability, as opposed to bounded deviation, is desirable in this application due to the impact on fuel consumption. From an empirical standpoint, the behavior of the SKMPC under uncertainty is demonstrated through numerical simulation in Section V.

D. Sequential Linearization Scheme

To reduce the prediction error introduced by the linearization in (12d), we employ a sequential linearization scheme [7], [21], where problem (12) is solved sequentially, each time re-linearizing the dynamics.

At each iteration of the sequential linearization, we update $\bar{\mathbf{X}}$ and $\bar{\mathbf{U}}$ with the optimal solution \mathbf{X}^* and \mathbf{U}^* to

problem (12) and recompute $\Phi_{j+1,j|k}$ and $\mathbf{c}_{j+1,j|k}$ using equations (3) and (13) along the updated the predicted states $\bar{\mathbf{x}}_{j|k}$ and pre-planned maneuvers $\bar{\mathbf{u}}_{j|k}$ for $j = 0, \dots, N-1$.

Algorithm 1 summarizes the SKMPC with the sequential linearization scheme. At t_k , the algorithm requires as input the control horizon $t_{0|k}, \dots, t_{N-1|k}$, the current state estimate $\hat{\mathbf{x}}_k^-$, the baseline state history $\mathbf{x}_{\text{ref}}(t)$, terminal constraint set $\mathcal{X}_{\text{targ}}(t_{N-1|k})$, admissible control set \mathcal{U} , and the maximum number of iterations for linearization M . In line 1, $\bar{\mathbf{X}}$ is initialized using the current state estimate $\hat{\mathbf{x}}_k^-$ and the baseline state at $t_{j|k}$ for $j = 1, \dots, N-1$, denoted by $\mathbf{x}_{\text{ref}}(t_{j|k})$. Note that solving (12) does not require a dynamically feasible initial sequence $\bar{\mathbf{x}}_{j|k}$. In line 4, `linearize` constructs the linearized expressions $\Phi_{j+1,j|k}$ and $\mathbf{c}_{j+1,j|k}$ around the a priori solution $\bar{\mathbf{x}}_{j|k}$ and $\bar{\mathbf{u}}_{j|k}$ for $j = 0, \dots, N-1$. In line 5, `SOC` solves (12) with a convex solver, e.g., here we use the interior point solver Clarabel [22]. The algorithm terminates once the solution \mathbf{X}^* and \mathbf{U}^* to (12) satisfies the nonlinear dynamics (4) for $j = 0, \dots, N-2$, and returns the earliest control $\mathbf{u}_{0|k}$ for execution. Then, the spacecraft remains in the corrected orbit until the next maneuver instance t_{k+1} , at which time Algorithm 1 is called again to compute $\mathbf{u}_{0|k+1}, \dots, \mathbf{u}_{N-1|k+1}$.

Algorithm 1 SKMPC with Sequential Linearization

Inputs: $t_{0|k}, \dots, t_{N-1|k}, \hat{\mathbf{x}}_k^-, \mathbf{x}_{\text{ref}}(t), \mathcal{X}_{\text{targ}}(t_{N-1|k}), \mathcal{U}, M$

- 1: $\bar{\mathbf{X}} \leftarrow [\hat{\mathbf{x}}_k^- \ \mathbf{x}_{\text{ref}}(t_{1|k}) \ \dots \ \mathbf{x}_{\text{ref}}(t_{N-1|k})]$
- 2: $\bar{\mathbf{U}} \leftarrow \mathbf{0}_{3 \times N}$
- 3: **while** true **do**
- 4: $\Phi, \mathbf{c} \leftarrow \text{linearize}(\bar{\mathbf{X}}, \bar{\mathbf{U}})$
- 5: $\mathbf{X}^*, \mathbf{U}^* \leftarrow \text{SOC}(\hat{\mathbf{x}}_k^-, \Phi, \mathbf{c}, \mathcal{X}_{\text{targ}}(t_{N-1|k}), \mathcal{U})$
- 6: **if** (4) is satisfied for $j = 0, \dots, N-2$ **then**
- 7: **break**
- 8: **end if**
- 9: $\bar{\mathbf{X}}, \bar{\mathbf{U}} \leftarrow \mathbf{X}^*, \mathbf{U}^*$ ▷ Update for next iteration
- 10: **end while**

Outputs: $\mathbf{u}_{0|k}^*$

IV. SIMULATION SETUP

The SKMPC is tested on a realistic SK scenario for a spacecraft flying on the NRHO. The simulation consists of recursively applying Algorithm 1 for an extended number of revolutions spanning multiple years, subject to navigation errors from the EKF, as well as modeling errors, control execution errors, and impulses due to momentum wheel desaturation.

A. Error Models

The filter is initialized with initial covariance $\mathbf{P}_{0|0} = \text{diag}([\sigma_{r_0}^2, \sigma_{r_0}^2, \sigma_{r_0}^2, \sigma_{v_0}^2, \sigma_{v_0}^2, \sigma_{v_0}^2])$ and initial state estimate

$$\hat{\mathbf{x}}_0^+ = \mathbf{x}(t_0) + \mathcal{N}(\mathbf{0}_{6 \times 1}, \mathbf{P}_{0|0}),$$

$$\mathbf{x}(t_0) = \mathbf{x}_{\text{ref}}(t_0) + \mathcal{N}(\mathbf{0}_{6 \times 1}, \text{block-diag}(\sigma_{r_0}^2 \mathbf{I}_3, \sigma_{v_0}^2 \mathbf{I}_3)),$$

where $\mathbf{x}(t_0)$ is the true state at the initial epoch t_0 , $\mathbf{x}_{\text{ref}}(t_0)$ is the baseline state at t_0 , σ_{r_0} is the initial position standard

deviation, and σ_{v_0} is the initial velocity standard deviation. At each revolution, when the spacecraft arrives at $\theta(t) = \theta_{\text{man}} = 200^\circ$, a maneuver is computed using the predicted state of the filter, \hat{x}_k^- and a control horizon defined by (11). The true state of the spacecraft is imparted with a corrupted maneuver using the Gates model [23]. We incorporate dynamics errors, which consist of variation in SRP magnitude and random impulses due to momentum wheel desaturation [3]. The former is modeled by relative perturbations $\delta(A/m)$ and δC_r on A/m and C_r in the \mathbf{a}_{SRP} term of (1), and the latter is modeled by an additive velocity perturbation $\Delta \mathbf{v}$ with random direction and magnitude when the spacecraft arrives at $\theta(t) = \theta_{\text{desat}}$, where θ_{desat} are desaturation true anomalies dictated by mission requirements [3].

Table I summarizes the error parameters, corresponding to the assumed levels of uncertainties for the Gateway [3], along with the selected process noise parameter σ_p . Note that the choice of σ_p is dependent on the canonical scales in which the dynamics are expressed.

TABLE I: Simulation parameters

Simulation parameter	Value
Average SRP A/m , m^2/kg	315/17900
Average SRP C_r	2
SRP rel. $\delta(A/m)$ 3- $\sigma_{A/m}$, %	30
SRP rel. δC_r 3- σ_{C_r} , %	15
Desaturation velocity magnitude 3- σ_{desat} , cm/s	1.0
	0° or
Desaturation true anomaly θ_{desat} , deg	$330^\circ, 0^\circ$ or $330^\circ, 0^\circ, 30^\circ$
Maneuver rel. magnitude error 3- $\sigma_{\mathbf{u},\text{rel}}$, %	1.5
Maneuver abs. magnitude error 3- $\sigma_{\mathbf{u},\text{abs}}$, mm/s	1.42
Maneuver execution direction error 3- $\sigma_{\Delta \mathbf{v},\text{dir}}$, deg	1°
Initial position standard deviation 3- σ_{r_0} , km	10
Initial velocity standard deviation 3- σ_{v_0} , mm/s	10
Range measurement 3- σ_r , m	1
Range-rate measurement 3- $\sigma_{\dot{r}}$, mm/s	0.1
Process noise parameter σ_p	5×10^{-5}

B. Navigation Update Model

In accordance with the typical operation of ground-based tracking, we assume measurements are provided during *tracking windows*, each lasting $\Delta t_{\text{track}} = 1$ hour. Let t_k and $t_{k+1} \approx t_k + T$, where $T \approx 6.55$ days is the approximate orbital period of the NRHO, denote two consecutive control epochs, such that $\theta(t_k) = \theta(t_{k+1}) = \theta_{\text{man}}$. In each revolution, we consider one post-maneuver tracking window starting 12 hours after t_k , and three pre-maneuver tracking windows, starting 72, 48, and 7 hours before t_{k+1} . During each tracking window, we provide $N_{\text{meas}} = 10$ equally spaced measurements. At t_{k+1} , Algorithm 1 uses the EKF's predicted state estimate following the latest measurement update provided $t_{k+1} - 7 + \Delta t_{\text{track}} = 6$ hours earlier.

C. Control Trigger Condition

To improve the delta-V performance of the SKMPC under navigation and execution errors, we consider a trigger condition to determine whether a maneuver is required. This

condition checks if the state propagated forward without control until t_N lies within an ellipsoid about the baseline with radii $\epsilon_{r,\text{trig}}$ in position and $\epsilon_{v,\text{trig}}$ in velocity components,

$\|\bar{\mathbf{r}}_{N-1|k} - \mathbf{r}_{f,\text{ref}}\|_2 \leq \epsilon_{r,\text{trig}}$, $\|\bar{\mathbf{v}}_{N-1|k} - \mathbf{v}_{f,\text{ref}}\|_2 \leq \epsilon_{v,\text{trig}}$, where $\bar{\mathbf{x}}(t_{N-1|k}) = [\bar{\mathbf{r}}_{N-1|k}^T \bar{\mathbf{v}}_{N-1|k}^T]^T$. Tolerances $\epsilon_{r,\text{trig}}$ and $\epsilon_{v,\text{trig}}$ do not need to be the same as ϵ_r and ϵ_v in (16). Choosing $\epsilon_{r/v} < \epsilon_{r/v,\text{trig}}$ makes the closed loop more robust against uncertainties; recursive feasibility is recovered by modifying the proof that is based on $\epsilon_{r,\text{trig}}$ and $\epsilon_{v,\text{trig}}$ instead of ϵ_r and ϵ_v , still assuming a sufficiently large u_{max} .

V. NUMERICAL RESULTS

We conduct a Monte-Carlo simulation, where each sample consists of navigating and performing SK over 300 revolutions along the baseline NRHO of the Gateway generated by NASA [1], corresponding to over 5.3 years. We compare the performance of SKMPC against XAC. With the SKMPC, we use $N_{\text{rev}} = 8$ and $u_{\text{max}} = 1 \text{ m/s}$, with triggering thresholds $\epsilon_{r,\text{trig}} = 100 \text{ km}$ and $\epsilon_{v,\text{trig}} = 20 \text{ m/s}$, and terminal constraint radii $\epsilon_r = 25 \text{ km}$ and $\epsilon_v = 5 \text{ m/s}$. With XAC, we use $N_{\text{rev}} = 7$, i.e., targeting the 7th perilune downstream, with $\epsilon_{v_x} = 1 \text{ m/s}$. From preliminary experiments, we note that we are able to use a larger N_{rev} with the SKMPC due to (1) the multiple-shooting formulation and (2) reduced sensitivity targeting conditions around apolune, compared to XAC, which is a single-shooting formulation that targets conditions around perilune. All thresholds are defined in \mathcal{F}_{EM} . We conduct three cases using 1, 2, and 3 desaturation events per revolution, at θ_{desat} provided in Table I.

The dynamics is integrated using the explicit Runge-Kutta Prince-Dormand (8,9) method from the GNU Scientific Library [24], rescaling (1) with canonical distance unit $\text{DU} = 10^5 \text{ km}$ and velocity unit $\text{VU} \triangleq \sqrt{\mu/\text{DU}}$ to improve numerical conditioning, with relative and absolute tolerances set to 10^{-14} . The SKMPC takes an average of 2.57 sec to solve on a single Intel i7-12700 CPU; the majority of the computational effort comes from propagating the STMs.

A. Navigation Performance

We first look at the navigation estimates provided to the controller. Figure 3 shows the estimation error of the EKF for the case involving 3 desaturation events and controlled by SKMPC; only the first 60 days are shown for the sake of clarity, as the filter performance is qualitatively similar across the remaining 240 days. The filter performance for 1 and 2 desaturation events, controlled by either the SKMPC or XAC, is qualitatively similar. For assessing SK activities, we focus on navigation performance at the maneuver time. Table II shows the numerical 3- σ pre-maneuver state estimation error with 1, 2, and 3 desaturation events using the SKMPC; errors using XAC are at similar orders of magnitude. As the number of desaturation events increases, the navigation error at the control epoch gets worse. The spikes in velocity estimate errors and corresponding covariance correspond to perilune passes, where the spacecraft undergoes a rapid change in direction of motion as it flies by the Moon.

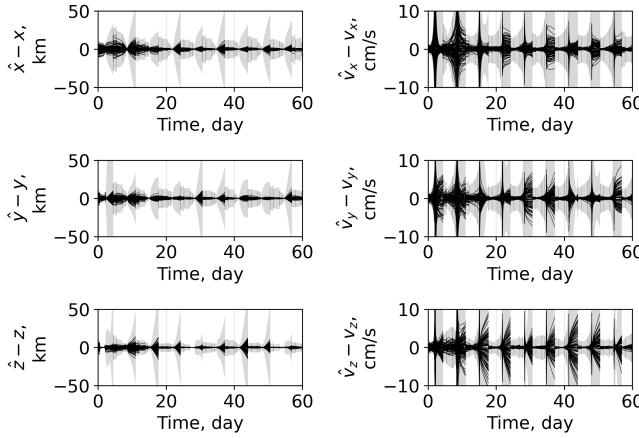


Fig. 3: Estimation error in \mathcal{F}_{Inr} with 3 desaturation events

B. Cost Performance

Tables III and IV provide cost statistics using the SKMPC and XAC. With both controllers, the SK cost increases as the navigation performance worsens; across varying numbers of desaturation events, the SKMPC provides both a lower yearly cost mean and standard deviation compared to XAC. The cumulative cost histories with 3 desaturation events per revolution are shown in Figures 4(a) and 4(b). With the SKMPC, the cumulative cost history follows a predominantly linear trend, indicating the controller is applying the appropriate level of control effort to keep the spacecraft motion near the baseline despite the uncertainties. In contrast, with XAC, the growth of the cumulative cost varies significantly based on the Monte-Carlo sample, which is a sign of the controller's susceptibility to uncertainties that lead to phase deviation.

TABLE II: Pre-maneuver state estimation error in \mathcal{F}_{EM}

Number of desaturation events	1	2	3
$(\hat{x} - x)$ 3- σ , km	0.924	1.041	1.128
$(\hat{y} - y)$ 3- σ , km	1.068	1.311	1.492
$(\hat{z} - z)$ 3- σ , km	0.635	0.677	0.711
$(\hat{v}_x - v_x)$ 3- σ , cm/s	0.213	0.222	0.228
$(\hat{v}_y - v_y)$ 3- σ , cm/s	0.700	0.927	1.086
$(\hat{v}_z - v_z)$ 3- σ , cm/s	0.101	0.119	0.133

TABLE III: Cost statistics using SKMPC

Number of desaturation events	1	2	3
Yearly mean, cm/s	109.96	153.48	186.83
Yearly standard deviation, cm/s	8.24	9.99	12.35
Yearly 95 th percentile, cm/s	123.21	169.34	208.63

TABLE IV: Cost statistics using *x*-axis crossing control

Number of desaturation events	1	2	3
Yearly mean, cm/s	111.94	135.76	155.29
Yearly standard deviation, cm/s	33.26	30.77	26.94
Yearly 95 th percentile, cm/s	178.78	192.55	198.66

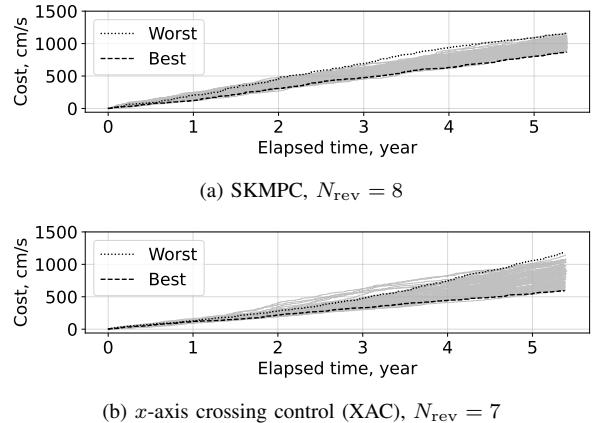


Fig. 4: Cumulative cost history with 3 desaturation events

C. Tracking Performance

We now analyze the tracking capability of the SKMPC and XAC. To isolate the effect of phase deviation, we compare the epochs and states at perilune passes of the controlled trajectory to the corresponding perilune passes of the baseline. Figures 5 and 6 show the deviation of the epoch and state in \mathcal{F}_{EM} at each perilune passage. Note that in \mathcal{F}_{EM} , perilunes occur approximately along the $+z$ axis, with the spacecraft's motion approximately perpendicular to the position vector; thus, the error is expected to be larger in z compared to x and y in position components, and v_x , and v_y compared to v_z in velocity components.

XAC experiences secular growth in phase deviations; the epoch deviation at perilune, while varying largely on the Monte-Carlo realization, reaches the order of several days, with passes deviating by up to 500 km in position and 200 m/s in velocity. By augmenting XAC with residual on perilune epoch deviation, Davis et al. [3] reports deviations of up to 80 km in position and 48 minutes in epoch.

In contrast, through its full-state tracking property, the SKMPC keeps the perilune epoch deviation to about 30 minutes, with position deviations of about 50 km and velocity deviations of about 10 m/s. In general, tighter tracking of the baseline is desirable since more stringent requirements can be met with regard to the spacecraft design or payload operations that require remaining closer to the intended path. For instance, the NRHO baseline for the Gateway is designed to be free of any Earth-shadowing eclipses [2], and tight tracking despite the dynamical uncertainties and errors can ensure that no such eclipse occurs during the flight.

VI. CONCLUSION

We proposed an economic MPC for the SK problem on the NRHO. This SKMPC achieves full-state tracking by taking into account multiple maneuvers within its control horizon. The controls are placed one revolution apart, making our approach operationally compliant to the single maneuver-per-revolution requirement typical in space missions on the NRHO. Through full-state tracking, the SKMPC overcomes

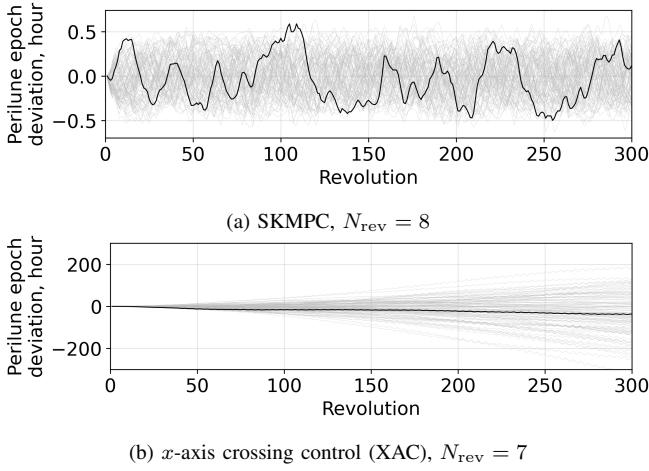


Fig. 5: Epoch deviation at perilune passes

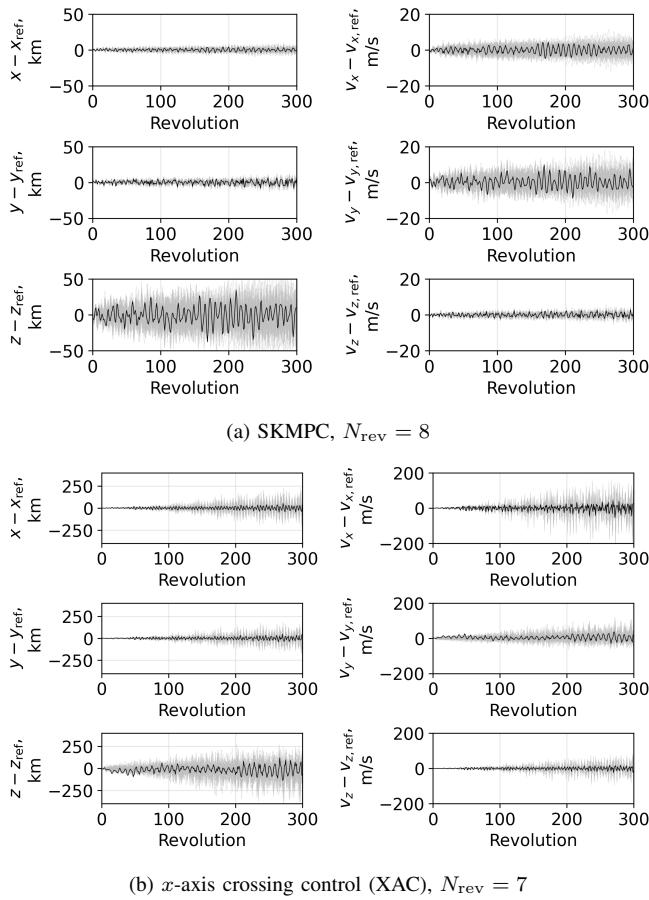


Fig. 6: State deviation in \mathcal{F}_{EM} at perilune passes

the issue of phase deviation encountered by other state-of-the-art SK schemes with a single maneuver per revolution. We tested the SKMPC with output-feedback using an EKF with range and range-rate measurements, and subject to realistic error models. Our approach achieves cumulative maneuver costs comparable to SK approaches proposed in the astrodynamics literature, while resulting in tighter tracking of the baseline orbit in both space and phase without additional ad-hoc heuristics, as required in XAC.

REFERENCES

- [1] D. E. Lee, "Gateway Destination Orbit Model: A Continuous 15 Year NRHO Reference Trajectory," NASA, Tech. Rep., 2019.
- [2] E. M. Zimovan-Spreen, K. C. Howell, and D. C. Davis, "Dynamical Structures Nearby NRHOs with Applications to Transfer Design in Cislunar Space," *Journal of the Astronautical Sciences*, vol. 69, no. 3, pp. 718–744, jun 2022.
- [3] D. C. Davis, S. T. Scheuerle, D. A. Williams, F. S. Miguel, E. M. Zimovan-Spreen, and K. C. Howell, "Orbit Maintenance Burn Details for Spacecraft in a Near Rectilinear Halo Orbit," in *AAS/AIAA Astrodynamics Specialists Conference*, 2022.
- [4] B. Cheetham, T. Gardner, A. Fornas, E. Kayser, and M. Clarkson, "CAPSTONE: A Unique CubeSat Platform for a Navigation Demonstration in Cislunar Space," in *ASCEND 2022*. Reston, Virginia: American Institute of Aeronautics and Astronautics, 2022, pp. 1–10.
- [5] K. C. Howell, "Three-dimensional, periodic, 'halo' orbits," *Celestial Mechanics*, vol. 32, no. 1, pp. 53–71, jan 1984. [Online]. Available: <http://link.springer.com/10.1007/BF01358403>
- [6] D. A. P. Williams, K. C. Howell, and D. C. Davis, "A Comparison of Station-Keeping Strategies for Halo Orbits," in *AAS/AIAA Astrodynamics Specialist Conference*, vol. 231, 2023, pp. 1–20.
- [7] Y. Shimane, K. Ho, and A. Weiss, "Optimization-Based Phase-Constrained Station-Keeping Control on Libration Point Orbit," in *AAS/AIAA Astrodynamics Specialists Conference*, 2024, pp. 1–19.
- [8] G. Misra, H. Peng, and X. Bai, "Halo orbit station-keeping using nonlinear MPC and polynomial optimization," in *AAS/AIAA Space Flight Mech. Meeting*, 2018.
- [9] P. Elango, S. Di Cairano, U. Kalabic, and A. Weiss, "Local Eigen-motion Control for Near Rectilinear Halo Orbits," in *2022 American Control Conference (ACC)*, 2022, pp. 1822–1827.
- [10] R. Quartullo, A. Garulli, and I. Kolmanovsky, "Periodic Model Predictive Control for Tracking Halo Orbits in the Elliptic Restricted Three-Body Problem," *IEEE Transactions on Control Systems Technology*, vol. 31, no. 5, pp. 1971–1981, 2023.
- [11] R. Padhi, A. Banerjee, S. Mathavaraj, and V. Srianish, "Computational Guidance Using Model Predictive Static Programming for Challenging Space Missions: An Introductory Tutorial with Example Scenarios," *IEEE Control Systems*, vol. 44, no. 2, pp. 55–80, 2024.
- [12] M. Shirobokov, S. Trofimov, and M. Ovchinnikov, "Survey of station-keeping techniques for libration point orbits," *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 5, pp. 1085–1105, 2017.
- [13] J. B. Rawlings, D. Angeli, and C. N. Bates, "Fundamentals of economic model predictive control," *Proceedings of the IEEE Conference on Decision and Control*, pp. 3851–3861, 2012.
- [14] D. Angeli, "Economic Model Predictive Control," in *Encyclopedia of Systems and Control*, J. Baillieul and T. Samad, Eds. London: Springer London, 2015.
- [15] Y. Shimane, S. D. Cairano, K. Ho, and A. Weiss, "Station-keeping on near-rectilinear halo orbits via full-state targeting model predictive control," in *2025 American Control Conference (ACC)*, 2025, pp. 1707–1712.
- [16] C. Acton, N. Bachman, B. Semenov, and E. Wright, "A look towards the future in the handling of space science mission geometry," *Planetary and Space Science*, vol. 150, no. January 2017, pp. 9–12, 2018.
- [17] W. McClain and D. Vallado, *Fundamentals of Astrodynamics and Applications*, ser. Space Technology Library. Springer Netherlands, 2001.
- [18] S. Särkkä, *Bayesian Filtering and Smoothing*, ser. Institute of Mathematical Statistics Textbooks. Cambridge University Press, 2013.
- [19] J. R. Carpenter and C. N. D'souza, "Navigation Filter Best Practices," NASA Engineering and Safety Center, Tech. Rep. TP-2018-219822 Navigation, 2018.
- [20] J. Rawlings, D. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, 2017. [Online]. Available: <https://books.google.com/books?id=MrJctAEACAAJ>
- [21] P. Elango, S. Di Cairano, K. Bernorp, and A. Weiss, "Sequential linearization-based station keeping with optical navigation for NRHO," in *AAS/AIAA Astrodynamics Specialist Conference*, 2022.
- [22] P. J. Goulart and Y. Chen, "Clarabel: An interior-point solver for conic programs with quadratic objectives," 2024.
- [23] C. R. Gates, "A Simplified Model of Midcourse Maneuver Execution Errors," Jet Propulsion Laboratory, Tech. Rep., 1963.
- [24] B. Gough, *GNU scientific library reference manual*. Network Theory Ltd., 2009.