(Preprint) AAS 25-664

OPTIMAL IMPULSIVE CONTROL OF CISLUNAR RELATIVE MOTION USING REACHABLE SET THEORY

Matthew Hunter[*], **Walter J. Manuel**[†], **and Simone D'Amico**[‡]

This work presents the first application of the state-of-the-art Koenig-D'Amico reachable set theory solver to cislunar, chaotic relative motion in the Circular-Restricted Three-Body Problem (CR3BP). The relative motion dynamics of two spacecraft, a chief and a deputy, in the CR3BP are formulated as a Linear Time-Variant (LTV) system, allowing the solver to find an optimal impulsive control maneuver plan. This methodology demonstrates robust and accurate control performance for both small and large reconfigurations over different CR3BP orbits and control windows. These capabilities are enhanced by a Model Predictive Control (MPC) architecture to reject all sources of control, navigation, and dynamic error. The performance of the proposed approach is validated by unit testing, Monte Carlo simulations, and comparisons to baseline models for spacecraft relative motion. Overall, this work demonstrates an optimal control methodology with the computational efficiency to be used on-board spacecraft, enabling the safe, effective, and efficient operation of Distributed Space Systems in cislunar space.

INTRODUCTION

The next era of space exploration is poised to be dually defined by activity in cislunar space and advancements in Distributed Space Systems (DSS). The race to return to the moon is currently spearheaded by NASA's Artemis Program, which aims to land humans on the lunar surface by 2027.[1] NASA is also partnering with international allies, including the European Space Agency (ESA), Japan Aerospace Exploration Agency (JAXA), and Canadian Space Agency (CSA), to develop the Lunar Gateway station, which will be placed in a Near-Rectilinear Halo Orbit (NRHO).[2] Additionally, the defense sector, specifically the US Space Force, has identified increasing Space Domain Awareness (SDA) in cislunar space as a priority.[3] This interest in cislunar space by government-focused sectors has already begun to drive investment from the commercial sector as well to help support and supplement civil and defense objectives.

Meanwhile, in Earth orbit, DSS are becoming increasingly commonplace, particularly in applications involving Rendezvous, Proximity Operations, and Docking (RPOD). DSS are defined as two or more spacecraft that work cooperatively, whether in formations, constellations, swarms, or other networks, to accomplish mission objectives not feasible with a single monolithic spacecraft. Multi-spacecraft systems will be an essential linchpin of the nascent cislunar ecosystem, as extensive RPOD operations of both crewed and uncrewed spacecraft are necessary to support lunar exploration and in-space logistics.

[*]PhD Candidate, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305.**
[†]PhD Candidate, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305.**
[‡]Associate Professor, Department of Aeronautics and Astronautics, Stanford University, Stanford, CA 94305.
**The first two authors have contributed equally to the paper.

However, the cislunar regime features more complicated and chaotic dynamics compared to Earth orbit, motivating further work to ensure the safe, effective, and efficient operation of cislunar DSS. This requires innovation both in the dynamics modeling and optimal control strategies. Dynamics modeling of two-body relative motion is fairly well explored, with widely-known approaches such as the Hill-Clohessy-Wiltshire equations,[4] the Yamanaka-Ankersen state transition matrices,[5] and the D'Amico Relative Orbital Elements (ROE).[6] However, many of the assumptions used to create those models do not always hold in the cislunar environment, necessitating different dynamics models for three-body relative motion. Recently, relative dynamics of the Circular Restricted Three-Body Problem (CR3BP) and the Elliptic Restricted Three-Body Problem (ER3BP) have been derived in the Local-Vertical-Local-Horizontal frame,[7,8] as well as in the velocity-oriented TNW frame[9] by linearizing around the chief spacecraft's position. Some authors have applied Linear Quadratic Regulator (LQR) control to a relative dynamics model obtained by linearizing the dynamics relative to a libration point, although this assumption potentially reduces accuracy in certain mission scenarios where the linearization is too coarse.[10] One optimal impulsive control method from Guffanti and D'Amico that leverages reachable set theory and primer vector theory,[11] has been applied to three-body absolute motion (linearized relative to a libration point),[12] and three-body relative motion.[13] However, these strategies all mostly focus on operational constraints, rather than provable optimality, and few offer the computational efficiency required to be implemented on-board a spacecraft.

Therefore, this work proposes a novel approach for modeling and controlling CR3BP relative spacecraft dynamics on-board over typical DSS control windows, resulting in three major contributions to the state of the art.

1. Two different methodologies are presented that leverage a linear instantaneous CR3BP relative dynamics model to generate a state transition matrix. The state transition matrix can be produced by either using a numerically efficient linear time-invariant approximation, breaking up the linear time-variant time horizon into smaller linear time invariant intervals, or using numerical integration, trading computational effort for dynamic accuracy. This facilitates the creation of a linear time-variant model of the relative CR3BP dynamics that propagates initial state and control actions to the end of a given time horizon.

2. This modeling approach enables the first application of the state of the art Koenig-D'Amico solver, or "KD solver," to the chaotic CR3BP dynamics of relative motion in the cislunar regime. The KD solver is a computationally-efficient, provably-optimal impulsive maneuver planner built upon a set of optimality conditions found through reachable set theory.[14] The KD solver has previously been applied to optimal path planning for satellites subject to two-body relative motion,[15] and the same reachable set optimality conditions have been extended to derive closed form solutions for DSS.[16]

3. The differing approaches for CR3BP state transition matrix creation are demonstrated in extensive simulation analysis. Variations in solver optimality, computational effort, and dynamic accuracy are quantitatively evaluated both in isolation and comparatively through Monte Carlo analysis against well known two-body state transition matrices from the literature. A simple Model Predictive Control setup is also implemented to demonstrate dynamic error rejection under typical operational uncertainties.

BACKGROUND

State Representation and Dynamics

This section provides technical background on modeling the CR3BP dynamics of absolute and relative spacecraft motion. This paper considers a system comprised of two spacecraft, a chief and a deputy, operating in the vicinity of the Moon and influenced by the gravitational forces of both the Earth and the Moon. As a result, the CR3BP is an appropriate dynamics model to characterize the motion of the two spacecraft.

In the CR3BP, there are two bodies of significant mass, such as the Earth and the Moon, accompanied by a spacecraft of relatively negligible mass.[17] The two larger bodies, known as the primaries, are modeled as point masses and rotate in circular orbits about their common center of mass. It is often convenient in CR3BP dynamic analysis to nondimensionalize the units of distance and time in the CR3BP and express motion in a synodic rotating frame that has the same inertial angular velocity as the primaries.

The system of deputy and chief spacecraft is visualized in Figure 1, along with the two key reference frames employed in this paper. The chief spacecraft is depicted as being on a Near-Rectilinear Halo Orbit (NRHO). NRHOs are a subclass of halo orbits, three-dimensional periodic orbits that naturally form about the equilibrium points of the CR3BP system. Close proximity to the Moon and favorable stability properties distinguish NRHOs from other halo orbits.[18]

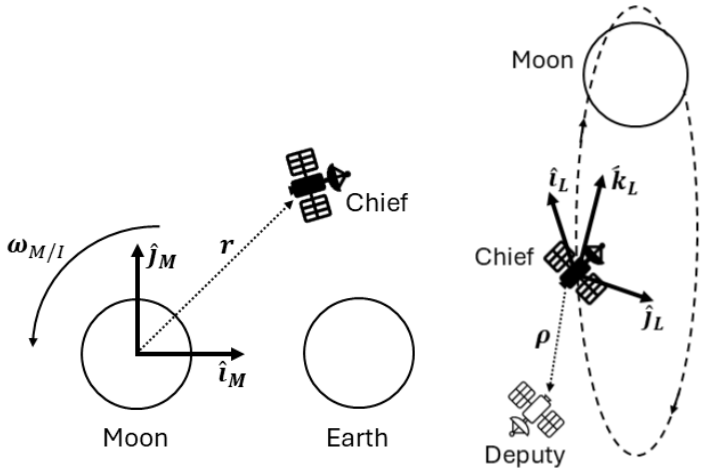

(a) Moon reference frame (b) LVLH reference frame

Figure 1: Key reference frames, vectors, and quantities used in this work. The chief spacecraft is depicted as being on a Near-Rectilinear Halo Orbit (NRHO).

First, the moon-centered synodic rotating frame $\mathcal{M}$, as shown in Figure 1a, is defined by unit vectors $\hat{\boldsymbol{i}}_M$, $\hat{\boldsymbol{j}}_M$, and $\hat{\boldsymbol{k}}_M$, where $\hat{\boldsymbol{i}}_M$ points from the moon to the Earth, $\hat{\boldsymbol{k}}_M$ is aligned with the angular momentum vector of the Earth-Moon system, and $\hat{\boldsymbol{j}}_M$ completes the right-handed triad. Second, a Local-Vertical Local Horizontal (LVLH) frame $\mathcal{L}$, centered on the chief spacecraft and anchored to the moon, is also employed as shown in Figure 1b. Frame $\mathcal{L}$ is defined by unit vectors $\hat{\boldsymbol{i}}_L$, $\hat{\boldsymbol{j}}_L$, and $\hat{\boldsymbol{k}}_L$, where $\hat{\boldsymbol{k}}_L$ points from the chief to the Moon, $\hat{\boldsymbol{j}}_L$ is anti-parallel to the angular momentum vector of the chief, and $\hat{\boldsymbol{i}}_L$ completes the right-handed triad. The LVLH frame is sometimes alternatively

defined as the Radial/Tangential/Normal (RTN) frame, using a permutation of the same unit vectors, $\{-\hat{\boldsymbol{k}}_L, \hat{\boldsymbol{i}}_L, -\hat{\boldsymbol{j}}_L\}$. The absolute position of the chief spacecraft relative to the Moon's center of mass is denoted as $\boldsymbol{r}$, and the relative position of the deputy spacecraft with respect to the chief is expressed as $\boldsymbol{\rho}$.

By linearizing the gravitational accelerations exerted on the deputy about the chief's position in the LVLH frame, the relative dynamics in the CR3BP system can be formulated as a Linear Time-Variant (LTV) system of equations in state space form as

$$\dot{\boldsymbol{x}} = \boldsymbol{A}(t)\boldsymbol{x} + \boldsymbol{B}\boldsymbol{u}, \tag{1}$$

where the state vector $\boldsymbol{x}$ is the Cartesian position and velocity of the deputy spacecraft in the LVLH frame, $\begin{bmatrix}\boldsymbol{\rho} & \dot{\boldsymbol{\rho}}\end{bmatrix}^\top$.[7] The plant matrix $\boldsymbol{A}(t)$ is defined as

$$\boldsymbol{A}(t) = \begin{bmatrix} \boldsymbol{0}_3 & \boldsymbol{I}_3 \\ \boldsymbol{A}_{\dot{\boldsymbol{\rho}}\boldsymbol{\rho}}(t) & -2\boldsymbol{\Omega}_{L/I}(t) \end{bmatrix}. \tag{2}$$

where $\boldsymbol{0}_3$ and $\boldsymbol{I}_3$ are 3x3 zero and identity matrices. The matrix $\boldsymbol{A}_{\dot{\boldsymbol{\rho}}\boldsymbol{\rho}}$ concisely combines the terms in the lower left quadrant, defined as

$$\begin{aligned} \boldsymbol{A}_{\dot{\boldsymbol{\rho}}\boldsymbol{\rho}}(t) = &- \left[\dot{\boldsymbol{\Omega}}_{L/I}\right]^{\mathcal{L}} - \boldsymbol{\Omega}_{L/I}^2(t) - \frac{\mu}{r^3}\left(\boldsymbol{I} - 3\frac{\boldsymbol{r}\boldsymbol{r}^\top}{r^2}\right) \\ &- \frac{1-\mu}{\|\boldsymbol{r}+\boldsymbol{r}_{em}\|^3}\left(\boldsymbol{I} - 3\frac{(\boldsymbol{r}+\boldsymbol{r}_{em})(\boldsymbol{r}+\boldsymbol{r}_{em})^\top}{\|\boldsymbol{r}+\boldsymbol{r}_{em}\|^2}\right), \end{aligned} \tag{3}$$

where $\boldsymbol{r}_{em}$ is the position vector of the earth relative to the moon, and μ is the mass parameter of the CR3BP system.[17] The matrix $\boldsymbol{\Omega}_{L/I} \in \mathbb{R}^{3\times3}$ is a skew-symmetric matrix associated with the angular velocity vector of frame $\mathcal{L}$ relative to the inertial frame $\mathcal{I}$, $\boldsymbol{\omega}_{L/I}$. Similarly, the skew-symmetric matrix $[\dot{\boldsymbol{\Omega}}_{L/I}]^{\mathcal{H}} \in \mathbb{R}^{3\times3}$ associated with the angular acceleration vector $[\dot{\boldsymbol{\omega}}_{L/I}]^{\mathcal{L}}$ is also introduced. These two matrices are defined as

$$\boldsymbol{\Omega}_{L/I} = \begin{bmatrix} 0 & -\omega_{L/I}^z & \omega_{L/I}^y \\ \omega_{L/I}^z & 0 & -\omega_{L/I}^x \\ -\omega_{L/I}^y & \omega_{L/I}^x & 0 \end{bmatrix}, \quad [\dot{\boldsymbol{\Omega}}_{L/I}]^{\mathcal{H}} = \begin{bmatrix} 0 & -\dot{\omega}_{L/I}^z & \dot{\omega}_{L/I}^y \\ \dot{\omega}_{L/I}^z & 0 & -\dot{\omega}_{L/I}^x \\ -\dot{\omega}_{L/I}^y & \dot{\omega}_{L/I}^x & 0 \end{bmatrix}. \tag{4}$$

The control matrix $\boldsymbol{B}$ transforms control actions $\boldsymbol{u}$ into instantaneous changes in state $\boldsymbol{x}$. Control actions are specified here to isolate changes in relative velocity, given as

$$\boldsymbol{B} = \begin{bmatrix} \boldsymbol{0}_3 \\ \boldsymbol{I}_3 \end{bmatrix}. \tag{5}$$

It is worth noting here that this work will use this LTV state space model both as the foundational model leveraged to create the proposed approach and as the reference ground truth of the validation simulation in the *Results* section below.

Problem Definition

This paper addresses the impulsive or instantaneous control of linear time-varying dynamic systems, as defined by the following. Consider a generic LTV dynamics system, given as

$$\boldsymbol{x}_f = \boldsymbol{\Phi}(t_0)\boldsymbol{x}_0 + \int_{t_0}^{t_f} \boldsymbol{\Phi}(\tau)\boldsymbol{B}(\tau)\boldsymbol{u}(\tau)d\tau, \tag{6}$$

where $\boldsymbol{x}_0 \in \mathbb{R}^n$ is the initial state at time t_0, $\boldsymbol{x}_f \in \mathbb{R}^n$ is the final state at time t_f, $\boldsymbol{\Phi}(t) = \boldsymbol{\Phi}(t, t_f) \in \mathbb{R}^{n \times n}$ is the State Transition Matrix (STM) that propagates a given state from time t to t_f, and $\boldsymbol{B}(t) \in \mathbb{R}^{n \times m}$ is the control matrix that converts control actions $\boldsymbol{u}(t) \in \mathbb{R}^m$ to changes in state $\boldsymbol{x}$. For simplicity, the pseudostate $\boldsymbol{\omega} \in \mathbb{R}^n$, the dynamic invariants of the control problem, and control matrix $\boldsymbol{\Gamma}(t) \in \mathbb{R}^{n \times m}$, the translation from control actions to changes in final state, are defined as,

$$\boldsymbol{\omega} = \boldsymbol{x}_f - \boldsymbol{\Phi}(t_0)\boldsymbol{x}_0, \quad \boldsymbol{\Gamma}(t) = \boldsymbol{\Phi}(t)\boldsymbol{B}(t). \tag{7}$$

Therefore, the optimal control problem (OCP) considered by this paper is defined as

$$\text{minimize: } \int_{t_0}^{t_f} f(\boldsymbol{u}(\tau), \tau) d\tau \quad \text{subject to: } \boldsymbol{\omega} = \int_{t_0}^{t_f} \boldsymbol{\Gamma}(\tau)\boldsymbol{u}(\tau) d\tau. \tag{8}$$

Cost function f is norm-like, such that it satisfies three properties: 1) f is defined for all maneuvers occurring within $t_0 \leq t \leq t_f$, 2) all sublevel sets of $f(\boldsymbol{u}(t), t)$ are convex and compact, and 3) $f(\alpha \boldsymbol{u}(t), t) = \alpha f(\boldsymbol{u}(t), t)$ for any $\alpha \geq 0$. Examples of norm-like cost functions and their corresponding spacecraft propulsion constraints include the L2-norm $||\boldsymbol{u}||_2$, where the spacecraft aligns a single thruster with the desired maneuver direction, and the L1-norm $||\boldsymbol{u}||_1$, where the spacecraft has a fixed attitude and three pairs of thrusters mounted on opposite sides. It is worth noting that the choice of a single cost function, rather than the more-general piecewise definition in previous works, is made for simplicity and does not effect the methodology detailed in later sections.

This work focuses on impulsive control actions, or instantaneous changes in state, assuming that the time required to complete each maneuver is short with respect to the secular and long-term periodic effects present in $\boldsymbol{\Phi}(t)$. Given there are no restrictive magnitude constraints on $\boldsymbol{u}$, time-varying control actions can be expressed as a set of k impulses, and the unconstrained control problem considered by the remainder of this paper is given as

$$\text{minimize: } c = \sum_{j=1}^{k} f_j(\boldsymbol{u}_j, t_j) \quad \text{subject to: } \boldsymbol{\omega} = \sum_{j=1}^{k} \boldsymbol{\Gamma}(t_j)\boldsymbol{u}_j, \tag{9}$$

where all $t_0 \leq t_j \leq t_f$ and c is the minimum cost objective.

Reachable Set Theory

The control solver proposed by this work is built upon the optimality conditions found through reachable set theory, and this analysis technique is detailed below. Let $U(c, t)$ be the set of control actions $\boldsymbol{u}$ with a cost no greater than c at time t, defined as

$$U(c, t) = \{\boldsymbol{u} : f(\boldsymbol{u}, t) \leq c\}. \tag{10}$$

Let $S(c, t)$ be the set of all $\boldsymbol{\omega}$ that can be reached by a single maneuver $\boldsymbol{u}$ of cost no greater than c at time t, given as

$$S(c, t) = \{\boldsymbol{y} : \boldsymbol{y} = \boldsymbol{\Gamma}(t)\boldsymbol{u}, \boldsymbol{u} \in U(c, t)\}. \tag{11}$$

This set is not generally convex. Let T be the discontinuous set of times t for admissible control actions in a control window bounded by t_0 and t_f, given as

$$T = \{t : t_0 \leq t \leq t_f\}. \tag{12}$$

Then, the set of all $\boldsymbol{\omega}$ that can be reached by a single maneuver $\boldsymbol{u}$ of cost no greater than c at any any time in T is defined as

$$S(c,T) = \cup_{t\in T} S(c,t), \tag{13}$$

which is also generally non-convex. Finally, let $S^*(c,T)$ be the set of $\boldsymbol{\omega}$ that can be reached by $k \geq 1$ control actions of combined cost no greater than c at times in T, defined equivalently as

$$S^*(c,T) = \{\boldsymbol{z} : \boldsymbol{z} = \sum\nolimits_{j=1}^{k} \boldsymbol{\Gamma}(t_j)\boldsymbol{u}_j, t_j \in T, \boldsymbol{u}_j \in U(c_j,t_j), \sum\nolimits_{j=1}^{k} c_j = c\} \tag{14}$$

$$= \{\boldsymbol{z} : \boldsymbol{z} = \sum\nolimits_{j=1}^{k} \alpha_j \boldsymbol{y}_j, \boldsymbol{y}_j \in S(c,T), \alpha_j \geq 0, \sum\nolimits_{j=1}^{k} \alpha_j = 1\}. \tag{15}$$

This shows that $S^*(c,T)$ is a linear combination and convex hull of $\boldsymbol{\omega}$ in $S(c,T)$. Returning to the optimal control problem in Eq. (9), the set $S^*(c,T)$ in Eq. (15) completely describes the dynamics constraint such that the problem can be defined as

$$\text{minimize: } c \quad \text{subject to: } \boldsymbol{\omega} \in S^*(c,T). \tag{16}$$

Computationally-Efficient Control Solver

This work extends a computationally-efficient solver designed by Koenig and D'Amico, denoted in this work as the "KD solver," for CR3BP application.[14] The KD solver is built upon a set of optimality conditions for control found through the Second-Order Cone Program (SOCP) dual of Eq. (16), given as

$$\text{maximize: } \boldsymbol{\lambda}^T\boldsymbol{\omega} \quad \text{subject to: } \max_{t\in T} g_{U(1,t)}(\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}) \leq 1. \tag{17}$$

where contact function $g_{U(1,t)}(\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}) = \max_{\boldsymbol{u}\in U(1,t)} \boldsymbol{\lambda}^T\boldsymbol{\Gamma}(t)\boldsymbol{u}$. First, consider an optimal solution to Eq. (16) and (17) defined as c_{opt} and $\boldsymbol{\lambda}_{opt}$. c_{opt} is the minimum cost to reach $\boldsymbol{\omega}$ if and only if $\boldsymbol{\omega}$ is on the boundary of $S^*(c_{opt},T)$. This boundary is described by the supporting hyperplane of $S^*(c_{opt},T)$ defined perpendicular to $\boldsymbol{\lambda}_{opt}$, and the optimal control input profile may not be unique if $\boldsymbol{\omega}$ lies on an edge or point of the boundary of $S^*(c_{opt},T)$. Second, for any $\boldsymbol{\omega}$ on the boundary of $S^*(c_{opt},T)$, an optimal control input sequence must exist such that $k \leq n$. Therefore, such a sequence must exist for any $\boldsymbol{\omega}$ of no more maneuvers than the dimensionality of $\boldsymbol{\omega}$. Third, control actions should only occur when $g_{U(1,t)}(\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}_{opt})$ reaches its maximum value in the control window T. This condition is reflected in the constraint of Eq. (17). Fourth, control actions should be in the direction of $\arg g_{U(1,t)}(\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}_{opt})$. The third and fourth conditions can be evaluated for an arbitrary $\boldsymbol{\lambda}$ and time t for an L2-norm cost function through the Cauchy Schwartz inequality, given as

$$g_{U(1,t)}(\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}) = ||\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}||_2, \tag{18}$$

$$\arg g_{U(1,t)}(\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}) = \frac{\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}}{||\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}||_2}. \tag{19}$$

These optimality conditions are visualized in Figure 2.

The KD solver consists of three major components: Initialization, Refinement, and Control Input Extraction. Initialization uses the direction of $\boldsymbol{\omega}$ as a good initial guess of the solution $\boldsymbol{\lambda}_{opt}$ of Eq. (17) to produce a set of initial candidate maneuvers times T_{est}, defined as the times where Eq. (18)

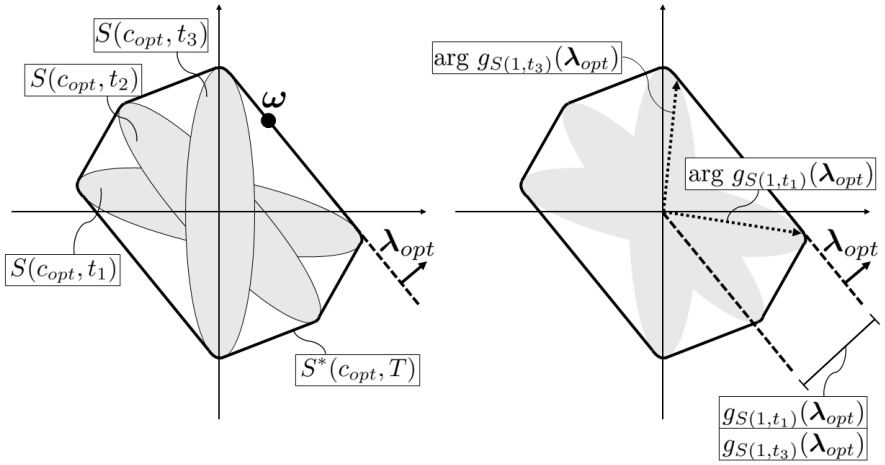

Figure 2: 2-dimensional reachable set formed as convex hull of states reached by single maneuvers at times $t_1 < t_2 < t_3$ (left). Optimality conditions evaluated on the boundary of $S^*(1, T)$ through contact functions (right).[15]

reaches a large value relative to other times in T. Refinement iteratively solves Eq. (17) to generate a new $\boldsymbol{\lambda}_{opt}$ for the times in T_{est}. This new $\boldsymbol{\lambda}_{opt}$ is used to remove sub-optimal candidate maneuver times from T_{est} whose contact function Eq. (18) falls below a threshold ϵ_{remove}, given as

$$g_{U(1,t)}(\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}_{opt}) < 1 - \epsilon_{remove}, \tag{20}$$

and add optimal candidate maneuvers times from T to T_{est} that exceed an optimality tolerance ϵ_{cost}, given as

$$g_{U(1,t)}(\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}_{opt}) > 1 + \epsilon_{cost}. \tag{21}$$

Refinement is repeated until $g_{U(1,t)}(\boldsymbol{\Gamma}^T(t)\boldsymbol{\lambda}_{opt}) < 1 + \epsilon_{cost}$ for all $t \in T_{est}$. Control Input Extraction finds an unconstrained optimal maneuver plan, under only dynamic constraints, at the optimal maneuver times in T_{est} by computing the optimal maneuver directions from Eq. (19) and finding the maneuver magnitudes that minimize the resulting pseudostate error of the maneuver plan through a Quadratic Program (QP). Limiting control inputs to times in T_{est} rather than T ensures that maneuvers only occur at optimal times without an explicit definition of optimality in the QP.

METHODOLOGY

To apply the KD solver to the CR3BP relative motion problem, a control matrix $\boldsymbol{\Gamma}$, comprised of an STM $\boldsymbol{\Phi}$ and control matrix $\boldsymbol{B}$, must be derived to form the dynamics constraint present in Eq. (9). The control matrix $\boldsymbol{B}$ is simply chosen as the same used in the CR3BP state space equations in Eq. (5). Two different strategies are proposed by this work to create an STM from the CR3BP dynamics in Eq. (2): matrix exponential and numerical integration. For matrix exponential, a Linear Time-Invariant (LTI) STM $\boldsymbol{\Phi}_{LTI}$ can be produced from a dynamic plant matrix $\boldsymbol{A}$, given as

$$\boldsymbol{\Phi}_{LTI}(t, t_f) = e^{\boldsymbol{A}(t)(t_f - t)}. \tag{22}$$

Consider approximating the LTV dynamics as a series of LTI dynamic systems. By discretizing the control window into an arbitrary number of time steps l, the LTV STM $\boldsymbol{\Phi}$ can be defined as a

multiplicative sequence of LTI STMs, given as

$$\boldsymbol{\Phi}(t) = \Pi_{i=1}^{l-1} \boldsymbol{\Phi}_{LTI}(t_i, t_{i+1}), \tag{23}$$

where $t_l = t_f$. The accuracy of this approximation depends on both the instantaneous variance of the true nonlinear dynamic system and the density of the time discretization, resulting in a tradeoff between computational efficiency and dynamic accuracy.

For numerical integration, the LTV STM is found by integrating the linear instantaneous dynamics in plant matrix $\boldsymbol{A}$ from a given time t to the end of the control window t_f. The time derivative of the STM $\boldsymbol{\Phi}$ is defined as

$$\frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{\Phi}(t) = \boldsymbol{A}(t)\boldsymbol{\Phi}(t), \qquad \boldsymbol{\Phi}(t = t_0) = \boldsymbol{I} \tag{24}$$

Therefore, the numerical integration problem resulting in the LTV STM is given as

$$\boldsymbol{\Phi}(t) = \int_t^{t_f} \boldsymbol{A}(\tau)\boldsymbol{\Phi}(t, \tau - t)d\tau. \tag{25}$$

The initial condition of this integration is the identity matrix, which logically follows from the definition of an STM over zero elapsed time. Numerical integration captures the full time-variance of the CR3BP relative dynamics over the control window.

In summary, the two different strategies proposed here to create an STM from the CR3BP dynamics will demonstrate a tradeoff between computational performance and dynamic accuracy. Using closed-form matrix exponential solutions enables the LTV system to be efficiently approximated as a series of LTI systems. However, this approximation will lose accuracy in a highly variable dynamic environment and may need a prohibitively dense time discretization to recover usable control solutions. In contrast, numerical integration completely characterizes the CR3BP relative dynamics within the accuracy of the linear model, but requires more substantial computational resources to achieve this performance. Figure 3 outlines the overall open and closed loop architectures employed in this work, illustrating how an STM is generated from the LTV system through either of the proposed methods, and then that STM is used by the KD solver to solve for the optimal maneuver plan.

RESULTS

The two proposed CR3BP STM creation approaches are thoroughly tested with the KD solver through multiple methods of validation. First, *Initial Validation* consists of two relevant test cases for CR3BP DSS control to demonstrate and compare the capabilities of the matrix exponential and numerical integration approaches. Then, *Validation Against Two-Body Dynamics Models* incorporates several commonly-used two-body STMs from the literature that have demonstrated a high level of accuracy for Earth orbit to show that the matrix exponential and numerical integration approaches outperform a two-body approximation of CR3BP motion. Next, *Monte Carlo Validation* expands this comparative analysis to a series of randomized reconfigurations, to show that the conclusions from the test case analysis remain consistent over a variety of solver inputs. Finally, *Model Predictive Control* incorporates the matrix exponential approach and KD solver into a simple robust control architecture, resolving the maneuver plan over an extended control window to reject dynamic modeling errors. For all validation simulations, control solutions are produced "on-board" using the STM creation approach and corresponding LTV dynamics model. These maneuver plans

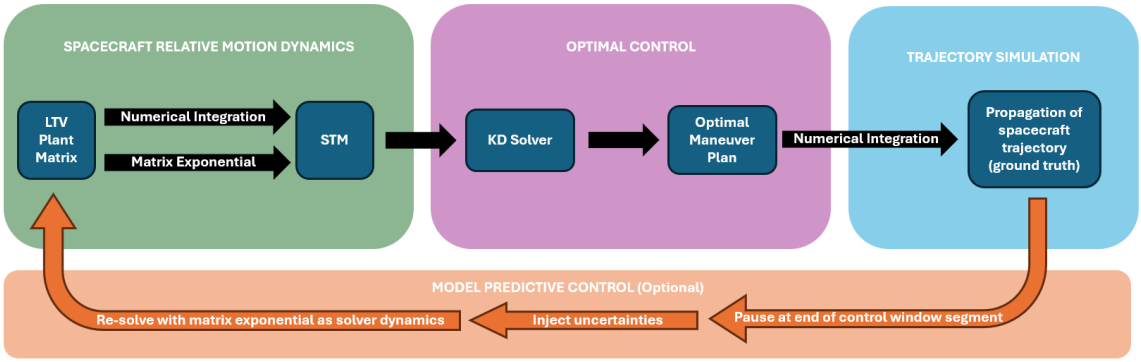

Figure 3: Block diagram depicting architecture of proposed methodology, from STM creation via two possible approaches, to solution of the OCP via the KD solver, to the implementation of that solution in numerically integrated ground truth dynamics. The option to close the loop with MPC control using the matrix exponential approach is also depicted. Other parameters not shown include the chief spacecraft absolute state, which is an input to form the LTV plant matrix; and inputs to the OCP inputs such as initial and final desired conditions, and control window length.

are then numerically propagated by the unperturbed equations of CR3BP motion, used as the ground truth. It is worth noting that this ground truth uses the same equations as that used to produce the matrix exponential and numerical integration STMs. However, inherent dynamic differences from the ground truth are produced in both cases. Ground truth errors in the matrix exponential approach are clearly created from the computationally efficient LTI approximation, while differing integration timesteps between the numerical integration approach and the simulation environment create small errors that are exaggerated by the chaotic dynamic environment around perilune. Maneuver plan cost, final state error, and solver runtime will be primarily used to quantify control performance under different STM creation approaches. All simulations are conducted on the same hardware, using a 3.80 GHz processor and 32GB RAM.

Initial Validation

The proposed solver and approaches to formulate the LTV control problem are first validated against two general reconfigurations in cislunar orbit. The initial absolute orbits of the chief spacecraft for both reconfigurations are given in the moon-centered synodic frame in Table 1.

Scenario	x (km)	y (km)	z (km)	$\dot{x}$ (km/s)	$\dot{y}$ (km/s)	$\dot{z}$ (km/s)
Reconfiguration 1	-13395	0	-70841	0	0.1055	0
Reconfiguration 2	-4909	29088	-14638	0.1080	-0.1647	0.4331

Table 1: Absolute states of the chief spacecraft at t_0 in the moon-centered synodic frame for initial validation reconfigurations.

Reconfiguration 1 represents a chief spacecraft in the 9:2 resonant NRHO at apolune, while Reconfiguration 2 represents a chief spacecraft in a 3:1 resonant halo orbit about to cross over perilune. The initial and final deputy relative states in LVLH frame for both reconfigurations are given in Table 2. The two initial reconfiguration scenarios were designed intentionally to represent plausible

Deputy State	x (km)	y (km)	z (km)	$\dot{x}$ (km/s)	$\dot{y}$ (km/s)	$\dot{z}$ (km/s)
Reconfiguration 1						
Initial	-300	-400	-200	0	0	0
Final	300	400	200	0	0	0
Reconfiguration 2						
Initial	-10	-0.3	-0.05	0	0	0
Final	0.1	0.3	0.05	0	0	0

Table 2: Initial and final deputy relative states in LVLH frame for initial validation reconfigurations.

hypothetical applications in cislunar space. Reconfiguration 1 changes the leader-follower order of the chief and deputy spacecraft, such that the deputy begins by "following" the chief with negative separation and ends by "leading" the chief with positive separation. Reconfiguration 1 has a 66.84h control window. Reconfiguration 2 involves a reduction in inter-spacecraft separation to conduct proximity operations and docking near the chief spacecraft and has a 33.52h control window. These reconfigurations will demonstrate the proposed solver's ability to accurately control CR3BP relative motion in real-world scenarios. An optimal control sequence for each reconfiguration is solved using both the matrix exponential and numerical integration STM approaches to compare resulting Δv costs, final state errors, and computational efficiency. These maneuver plans are simulated in a ground truth numerically integrated propagation of the unperturbed chief (nonlinear absolute CR3BP) and deputy (linear relative CR3BP) dynamics.

The KD solver is set up with a control domain T of 1001 candidate maneuver times for each reconfiguration, produced by uniformly splitting the control window into 1000 steps. The STM is precomputed at each candidate maneuver time before running the solver to isolate the computational effort associated with calculating the STMs from that required to find the optimal maneuver plan. The matrix exponential approach breaks the Reconfiguration 1 control window into 10min time steps and the Reconfiguration 2 control window into 20min time steps for the LTI approximation. The chosen cost function is the sum of L2-norms of RTN control actions at times in T, and control is admissible at all times in the control window. The initialization component of the KD solver is set to evaluate the contact function for every ten candidate times in T, taking the ten highest as the initial T_{est}, and the Control Input Extraction component uses an identity matrix to minimize the pseudostate error. The resulting control trajectories are displayed in Figure 4.

For both reconfigurations, the KD solver found almost identical maneuver plans for the matrix exponential and numerical integration STM approaches, illustrated by the overlay of trajectories and maneuvers. The chief's proximity to perilune plays a clear role in the resulting controlled and uncontrolled trajectories. The chief and deputy far from perilune in Reconfiguration 1 feature smooth, consistent changes in relative state with low curvature. Reconfiguration 2 demonstrates the highly variable dynamics near perilune with significant, high curvature changes in the deputy relative state path over a shorter control window than Reconfiguration 1. The contrasting advantages and disadvantages of the two STM approaches are well illustrated by the performance metrics associated with each reconfiguration in Table 3.

As expected, the similar trajectories and maneuvers directions in Figure 4 correspond to similar optimal maneuver plans costs found by both STM strategies for each reconfiguration. The STM and

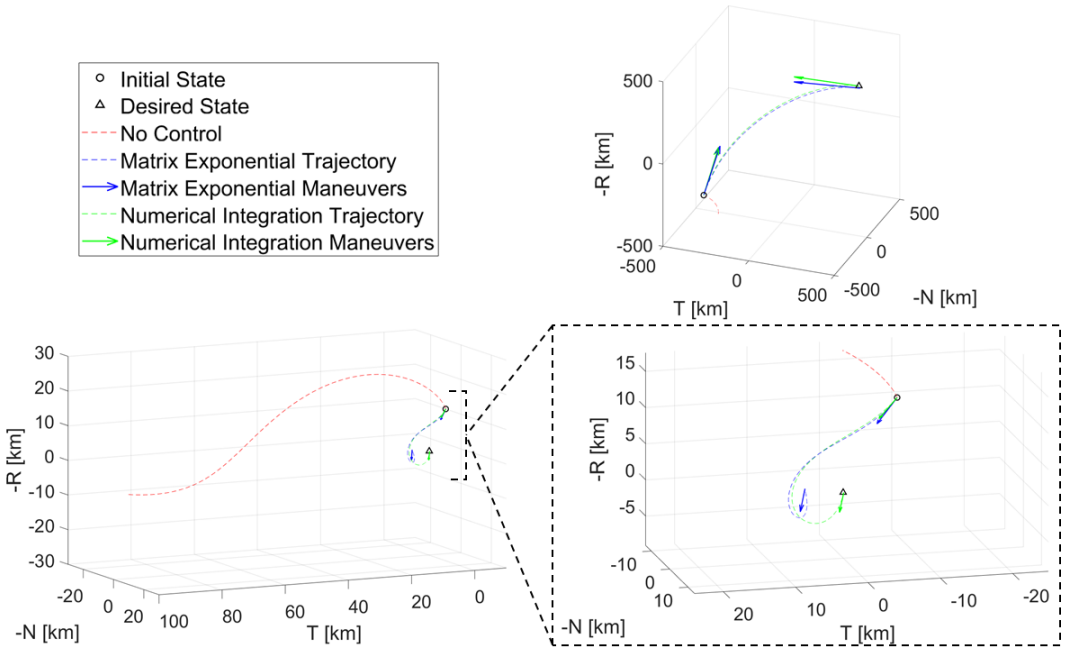

Figure 4: LVLH trajectories for Reconfiguration 1 (top row) and Reconfiguration 2 (bottom row).

STM Strategy	Cost (m/s)	RMS $\boldsymbol{x}_f$ Error (km)	STM Runtime (s)	Solver Runtime (s)
Reconfiguration 1				
Matrix Exp.	9.7644e-4	8.4613	0.55	0.96
Numerical Int.	9.9372e-4	0.8065	8.63	1.00
Reconfiguration 2				
Matrix Exp.	4.6352e-4	2.9950	13.36	26.84
Numerical Int.	4.8193e-4	0.0496	26.17	26.40

Table 3: Performance metrics for initial validation reconfigurations.

solver runtimes are both significantly longer for Reconfiguration 2 than Reconfiguration 1 despite having a far shorter control window. This is due to the proximity of perilune, causing dynamic variability that requires more numerical effort to both calculate matrix exponentials and numerically integrate, as well as preventing quick solver convergence. The computational advantages of the matrix exponential approach are clear here, taking only 6.37% and 51.05% of the numerical integration approach for the two reconfigurations. The two STM approaches also result in very similar solver runtimes, demonstrating that the choice of STM does not affect the computational performance of the optimization algorithm. In contrast, the dynamic accuracy advantages of the numerical integration approach are clear in the final state terminal control error, 9.53% and 1.66% of the matrix exponential approach for Reconfiguration 1 and 2, respectively. This difference in control accuracy can be further understood by examining the STM propagation errors resulting from the two STM approaches, shown in Figures 5 and 6 with y-axes on a log scale for ease of comparison.

The STM Root-Mean-Squared (RMS) propagation errors are found by propagating the initial

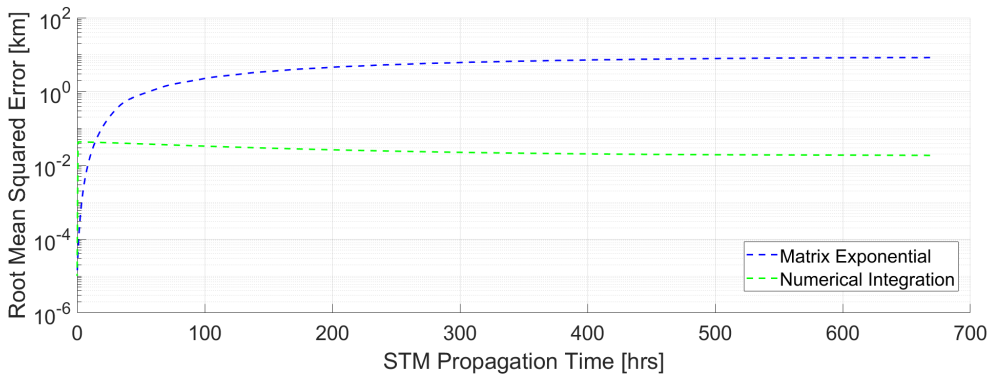

Figure 5: RMS propagation errors of each STM creation strategy from the initial chief and deputy states in Reconfiguration 1.

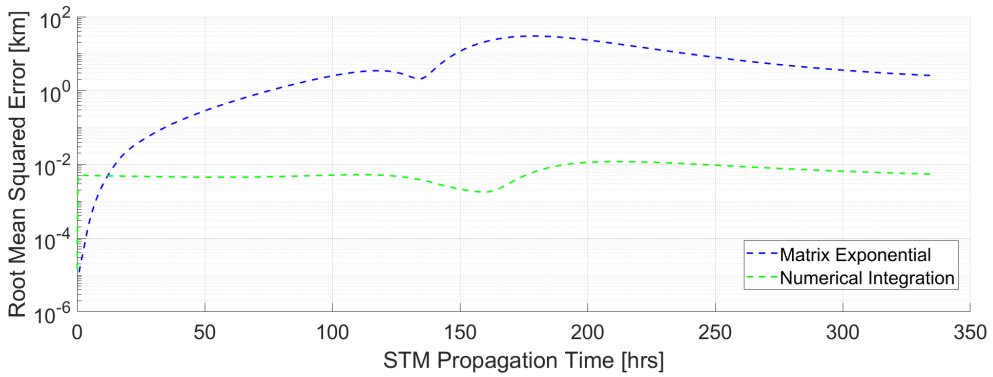

Figure 6: RMS propagation errors of each STM creation strategy from the initial chief and deputy states in Reconfiguration 2.

deputy state with the chosen STM and comparing the resulting state against the numerically integrated, unperturbed CR3BP ground truth over the same time interval. In both reconfigurations, the RMS propagation error of the numerical integration approach does not exceed 0.1km while the RMS error of the matrix exponential approach can reach or exceed 10km. This difference in relative accuracy matches that for the final state RMS error for both reconfigurations in Table 3. Again, the proximity of perilune has a clear effect on performance. Both STM approaches in Reconfiguration 1 reach a fairly steady-state RMS error for propagation times greater than 100h, while passing over perilune in Reconfiguration 2 results in an oscillation in RMS error that only approaches similar steady-state values to Reconfiguration 1 for propagation times greater than 300h. In general, the initial validation reconfigurations reflect the expected tradeoff in computational efficiency and dynamic accuracy present between the matrix exponential and numerical integration approaches. Additionally, the accurate linear dynamic model resulting from the numerical integration strategy both near and far from perilune demonstrates the solver's capability to produce accurate maneuvers plans in both consistent and highly variable dynamic environments.

Validation Against Two-Body Dynamics Models

The Hill-Clohessy-Wiltshire (HCW) and Yamanaka-Ankersen (YA) dynamics models are commonly used to accurately characterize and control the relative motion of DSS in two-body Earth orbit. These models are used here as a baseline to evaluate the performance of the two CR3BP LTV relative motion modeling approaches. The same problem setup from Reconfiguration 1 in the *Initial Validation* section was used for this comparative testing, with the exception of the control window, set to 167.1h to evaluate the performance over longer durations, and the timestep, reduced to 1 minute in an attempt to increase the accuracy of the two-body dynamics models. As before, the solved maneuver plans are then simulated in a ground truth numerical propagation of the chief and deputy CR3BP unperturbed dynamics.

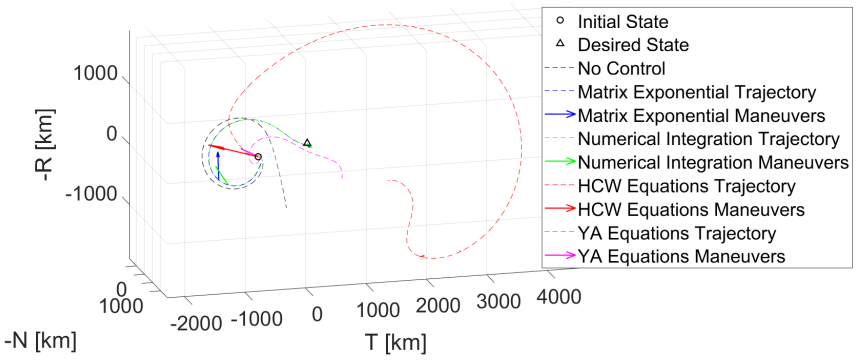

Figure 7: LVLH trajectories for Reconfiguration 1 with an extended control window.

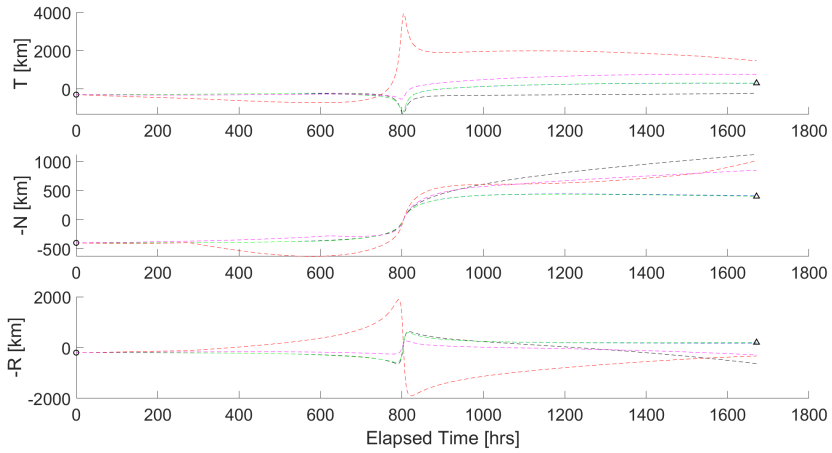

Figure 8: Time evolution of deputy LVLH position components for Reconfiguration 1 with an extended control window. Color coding is same as Figures 7 and 10.

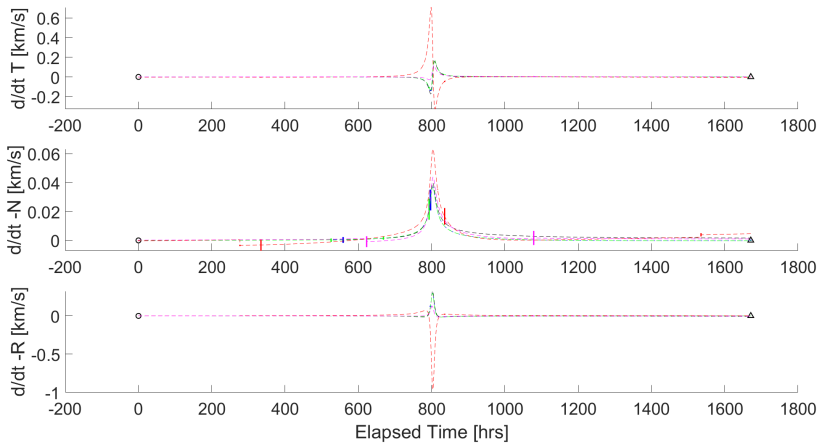

Figure 9: Time evolution of deputy LVLH velocity components for Reconfiguration 1 with an extended control window. Color coding is same as Figures 7 and 10.

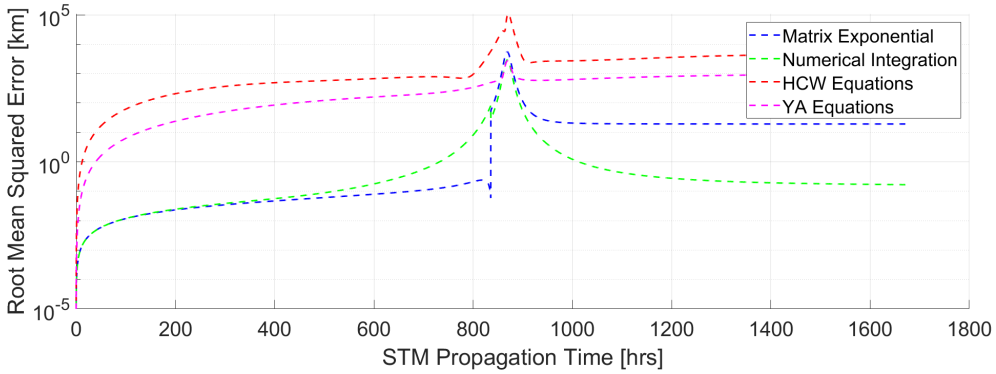

Figure 10: RMS propagation errors of each STM creation strategy from the initial chief and deputy states in Reconfiguration 1 with an extended control window.

Figures 7, 8, and 9 display the trajectories of each dynamics model, where the STMs generated by each of the five models result in five distinctly different maneuver plans when input to the KD solver. Figure 10 shows the final state propagation error of the different models, with the numerical integration approach yielding the lowest error, followed by the matrix exponential approach, while the two-body dynamics models all have drastically higher errors. A bifurcation between the CR3BP LTV and the two-body models can be seen almost immediately in the RMS error timeseries, and eventually this delta in error magnitude reaches over 100 km by the terminal state. Notably, all approaches show an increase in error when transiting perilune. Nevertheless, this example demonstrates the proposed CR3BP LTV models outperforming the baseline models, providing evidence of the benefits gained by employing dynamics models tailored to the CR3BP.

Monte Carlo Validation

A Monte Carlo simulation is the next step taken in robustly validating the CR3BP LTV approaches over a wide range of possible scenarios. The goal of the Monte Carlo simulation is to

1) Demonstrate clear performance above the baseline two-body models (HCW and YA), 2) Quantify the tradeoffs in accuracy and computational efficiency between the numerical integration and matrix exponential approaches, and 3) Evaluate the solver performance over a wide range of cislunar reconfiguration scenarios. The setup of the Monte Carlo simulation, including which parameters are varied, and what distributions are used for sampling, is described in Table 4. Deputy initial state, deputy final state, chief initial absolute Halo orbit, and chief phase along the Halo orbit are all sampled from a representative distribution of different separations and scenarios. The control window length is also chosen randomly to evaluate performance on short and long duration formation keeping or reconfiguration activities. In total, 100 test cases are sampled and used in the simulation. All five dynamics models are evaluated on the same 100 test cases for consistency. The time step used for the LTI approximation within the matrix exponential approach is set at 60s for all test cases, and the KD solver setup matches that used in *Initial Validation*.

Variable	Sampling Approach
Chief initial absolute state	Uniform sampling from 6000 known halo orbit states: 1000 states each from the 9:2, 4:1, 7:2, 3:1, 5:2, 2:1 halo orbits
Deputy initial relative position	Log-uniform distribution $\sim$ [-5000, 5000] km
Deputy initial relative velocity	Normal distribution, Mean = 0 km/s, Std. dev. = 0.001 km/s
Deputy final relative position	Log-uniform distribution $\sim$ [-5000, 5000] km
Deputy final relative velocity	Normal distribution, Mean = 0 km/s, Std. dev. = 0.001 km/s
Control window	Log-uniform distribution $\sim [0.1\pi, 4\pi]$ TU

Table 4: Variables of interest and associated sampling approaches for Monte Carlo simulation.

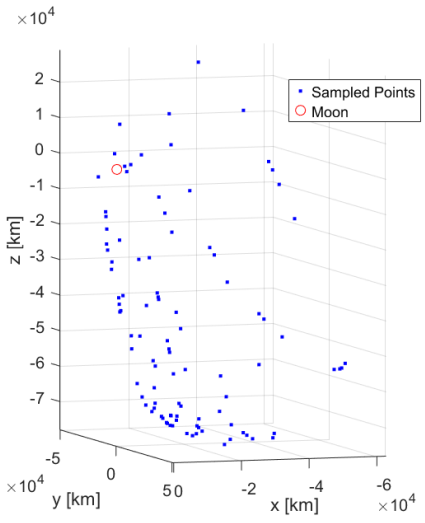

Figure 11: Chief spacecraft initial absolute states (in moon-centered synodic frame) used in Monte Carlo simulation, sampled randomly from 9:2, 4:1, 7:2, 3:1, 5:2, and 2:1 Halo orbits.

Figure 11 provides a visualization of the sampling approach taken for the chief spacecraft initial absolute state, drawing from multiple halo orbits and different phases along those orbits, including at perilune.

Table 5 reports the results of the Monte Carlo simulation. Clearly, at scale, the same trends

Metric	Matrix Exp.	Num. Int.	HCW	YA
Final Position Error (km) - Median	8.8225	9.8411	6208.8	3855.3
Final Position Error (km) - Min	7.2799e-2	7.2862e-2	26.469	27.665
Final Position Error (%) - Median	5.1399	3.5769	2620.9	1163.9
Final Position Error (%) - Min	7.4492e-3	2.9771e-3	17.707	9.7436
Optimal Cost (m/s) - Median	14.547	14.542	10.265	5.3085
Optimal Cost (m/s) - Mean	35.166	35.661	18.812	20.290
Optimal Cost (m/s) - Max	344.15	348.67	97.508	351.65
Optimal Cost (m/s) - Min	1.8811	1.4988	0.98025	0.57857
STM Runtime (s) - Median	1.9914	12.597	1.5156	1.8405
STM Runtime (s) - Mean	4.9494	22.369	3.4174	3.9759
STM Runtime (s) - Max	23.727	84.492	16.541	16.864
STM Runtime (s) - Min	0.19371	1.6645	0.13326	0.17465
Solver Runtime (s) - Median	7.06665	7.4127	6.7330	6.9115
Solver Runtime (s) - Mean	20.225	16.626	15.865	22.378
Solver Runtime (s) - Max	285.08	87.209	94.789	185.90
Solver Runtime (s) - Min	1.0622	1.0237	0.68229	0.98903

Table 5: Performance metrics for Monte Carlo simulation, comparing multiple dynamics models.

persist from the individual test case comparing the five dynamics models. The matrix exponential and numerical integration approaches had a median terminal position error of just 5.1% and 3.6%, respectively. In contrast, the HCW and YA models failed to model the dynamics accurately enough for the solver to produce any accurate solutions, with median terminal position error of over 1000%. When comparing the matrix exponential and numerical integration approaches, the data suggests that the tradeoffs in accuracy and efficiency favor the matrix exponential approach. The average time needed for STM generation is an order of magnitude lower for the matrix exponential approach compared to numerical integration, while the average solver runtimes are comparable. However, the matrix exponential approach does not suffer a significant drop in error compared to numerical integration. In short, the matrix exponential approach appears to be much faster than numerical integration, while only sacrificing a relatively small level of accuracy on average.

Model Predictive Control

The final component of analysis and validation for this work is to implement a Model Predictive Control (MPC) scheme. MPC rejects dynamic propagation errors, which stem from the inherent differences between the matrix exponential and the numerically integrated LTV model, and can also mitigate maneuver execution errors and navigation uncertainty. The naive MPC architecture begins with an initial control solution at the beginning of the control window and periodically produces a new maneuver plan for the remaining control window from the current state estimate. This strategy aims to leverage the more computationally efficient matrix exponential approach to enable accurate, robust closed loop control. For the MPC test case, the initial and desired final states of the deputy are given in Table 6, and the control window was 167.1h, to simulate a long-term reconfiguration. The MPC strategy divides this control window into 10 segments and then re-solves at the end of each segment, as the deputy spacecraft draws closer to the terminal state. The chief spacecraft initial absolute state is at apolune of the 9:2 NRHO. Typical maneuver execution and navigation errors

for small satellites in the CR3BP regime are considered as part of the MPC validation simulation environment. Navigation and maneuver execution performance are simulated by a set of zero-mean Gaussian distributions given in Table 7.

Deputy State	x (km)	y (km)	z (km)	$\dot{x}$ (km/s)	$\dot{y}$ (km/s)	$\dot{z}$ (km/s)
Initial	-3000	-4000	-2000	0	0	0
Final	3000	4000	2000	0	0	0

Table 6: Initial and final deputy relative states in LVLH frame for MPC reconfiguration.

Parameter	Covariance
Chief absolute position components	1^2 km^2
Chief absolute velocity components	0.01^2 km^2/s^2
Deputy relative position components	0.01^2 km^2
Deputy relative velocity components	0.001^2 km^2/s^2
Maneuver time	60^2 s^2
Maneuver magnitude	0.01^2 km^2
Maneuver direction	1^2 deg^2

Table 7: Navigation and maneuver execution uncertainties used in MPC reconfiguration.

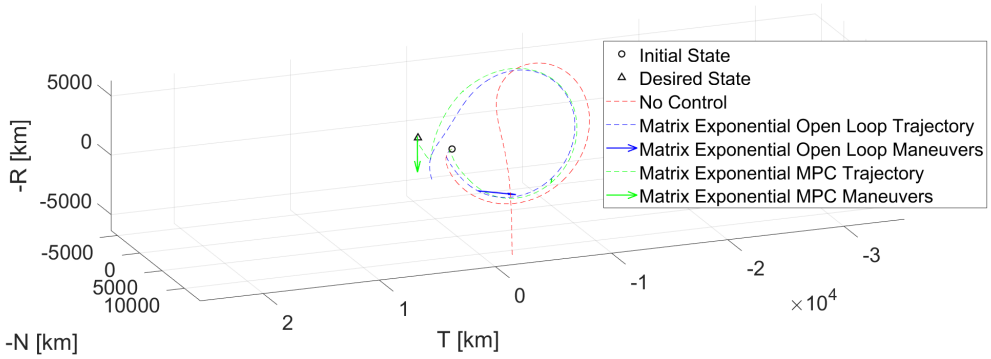

Figure 12: Trajectory of deputy in LVLH frame, comparing MPC to open loop control.

The results of the MPC simulation for a single test case can be visualized in Figures 12, 13, and 14. Open loop control and uncontrolled trajectories are included for comparison. Open loop control in this context is equivalent to the naive maneuver planning used in previous sections of this paper. The tradeoffs between MPC and open loop control are summarized in 8, and are consistent with expectations. The terminal position error from MPC is less than a tenth of that attained using open loop control, due to the MPC regularly re-planning as the time horizon shrinks. However, this increase in accuracy comes at the expense of Δv cost and runtime, where the metrics for MPC exceed those for open loop control by a factor of about four.

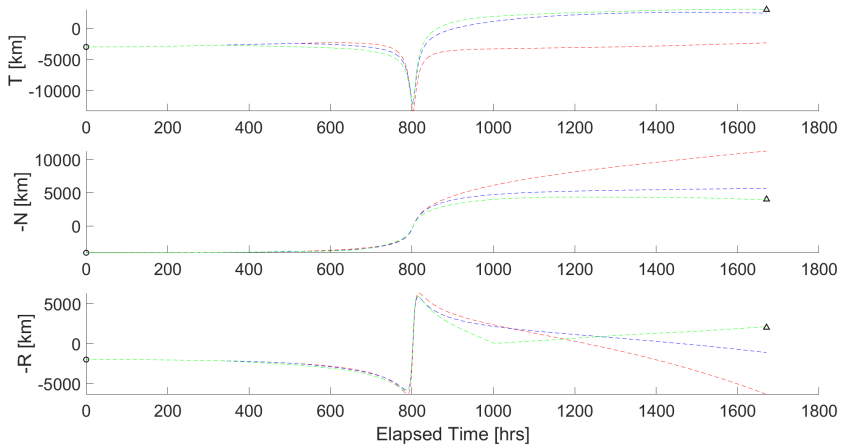

Figure 13: Time evolution of deputy position components in LVLH frame, comparing MPC to open loop control.

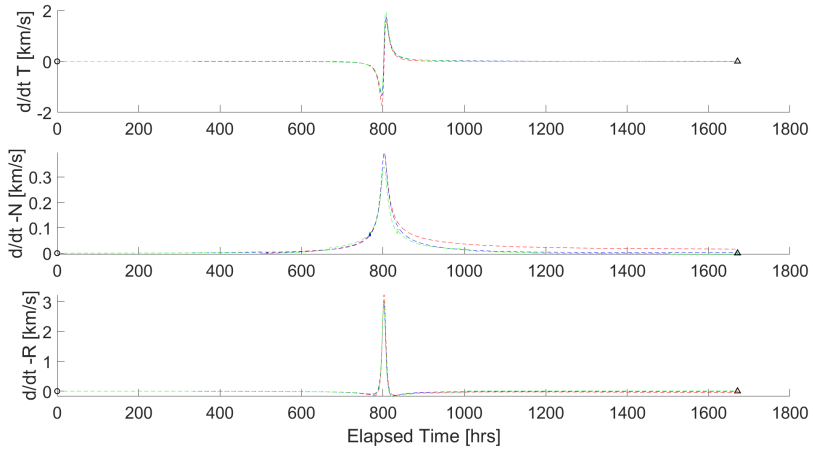

Figure 14: Time evolution of deputy velocity components in LVLH frame, comparing MPC to open loop control.

Metric	MPC	Open Loop
Terminal Position Error (km)	162.25	2087.0
Terminal Position Error (%)	3.0130	38.755
Optimal Cost (m/s)	92.885	23.131

Table 8: Performance metrics for MPC validation reconfigurations.

CONCLUSIONS

This work applies a generic reachable set theory optimal control approach for Linear Time-Variant (LTV) dynamic systems to cislunar relative spacecraft motion. First, a linear model of Circular Restricted Three-Body Problem (CR3BP) relative spacecraft motion is leveraged to create

an LTV dynamics model. The instantaneous changes in relative Cartesian state are numerically integrated to produce a State Transition Matrix (STM) that translates impulsive control actions from time of execution to the end of the control window. Second, the LTV model is incorporated into a reachable set theory-based solver that uses a set of optimality conditions to efficiently produce impulsive control profiles with provable optimality. Accommodations are made within the optimization algorithm to make it robust to the extreme variations in relative and absolute dynamics present near perilune. Third, this methodology demonstrates accurate control performance over short and long control windows, accomplishing both small and large reconfigurations over different CR3BP orbits. These capabilities are enhanced by a Model Predictive Control (MPC) architecture to reject all sources of control, navigation, and dynamic error.

The MPC results are promising, as they demonstrate the capability to handle uncertainties and perturbations through closed loop control. While the work presented in this paper is subject to CR3BP dynamics, it could be extended to higher fidelity models in the future. Adding perturbations such as SRP and lunar spherical harmonics to the ground truth simulation is a logical next step, with the ultimate goal of using a full-ephemeris model. Other higher fidelity models, such as the Elliptical Restricted Three-Body Problem (ER3BP) could be employed as well. It is important to note that since an LTV system is required to be compatible with the KD solver, there will always be a delta between the STM dynamics and the ground truth dynamics due to the necessary linearizations. Further investigation is required to determine whether the KD solver is capable of overcoming this delta in higher fidelity models, both in open loop and closed loop control approaches. Overall, this work demonstrates an optimal control methodology for cislunar applications with the computational efficiency to be used on-board spacecraft, enabling spacecraft autonomy beyond Earth orbit.

ACKNOWLEDGMENTS

This work is supported by Blue Origin (SPO #299266) as an Associate Member and Co-Founder of the Stanford's Center of AEroSpace Autonomy Research (CAESAR), and by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2146755. This article solely reflects the opinions and conclusions of its authors and not any of its sponsors.

REFERENCES

[1] M. Smith, D. Craig, N. Herrmann, E. Mahoney, J. Krezel, N. McIntyre, and K. Goodliff, "The artemis program: An overview of nasa's activities to return humans to the moon," *2020 IEEE aerospace conference*, IEEE, 2020, pp. 1–10.

[2] S. Fuller, E. Lehnhardt, C. Zaid, and K. Halloran, "Gateway program status and overview," *Journal of Space Safety Engineering*, Vol. 9, No. 4, 2022, pp. 625–628.

[3] M. J. Holzinger, C. C. Chow, and P. Garretson, "A primer on cislunar space," 2021.

[4] W. H. Clohessy and R. S. Wiltshire, "Terminal Guidance System for Satellite Rendezvous," *Journal of the Aerospace Sciences*, Vol. 27, Sept. 1960, pp. 653–658, 10.2514/8.8704.

[5] K. Yamanaka and F. Ankersen, "New State Transition Matrix for Relative Motion on an Arbitrary Elliptical Orbit," *Journal of Guidance, Control, and Dynamics*, Vol. 25, Jan. 2002, pp. 60–66, 10.2514/2.4875.

[6] S. D'Amico, *Autonomous Formation Flying in Low Earth Orbit*. PhD thesis, TU Delft, 01 2010.

[7] G. Franzini and M. Innocenti, "Relative Motion Dynamics in the Restricted Three-Body Problem," *Journal of Spacecraft and Rockets*, Vol. 56, Sept. 2019, pp. 1322–1337, 10.2514/1.A34390.

[8] F. Khoury and C. Lippe, "Relative Guidance, Navigation, and Control, in Multibody Gravitational Regimes," *AAS/AIAA Astrodynamics Specialist Conference*, Charlotte, NC, 2022.

[9] Y. Takubo, W. J. Manuel, and S. D'Amico, "Passively-Safe Optimal Cislunar Relative Motion using Local Toroidal Coordinates," *AAS/AIAA Spaceflight Mechanics Meeting*, Kauai, HI, 2025.

[10] J. S. Ardaens and S. D'Amico, "Control of Formation Flying Spacecraft at a Lagrange Point," *Technical Note*, 2008.

[11] T. Guffanti and S. D'Amico, "Integration Constants as State Variables for Optimal Path Planning," *2018 European Control Conference (ECC)*, IEEE, 2018, pp. 1–6.

[12] W. J. Manuel and S. D'Amico, "Optimal Control in the Circular Restricted Three-Body Problem using Integration Constants," *AAS/AIAA Astrodynamics Specialist Conference*, Charlotte, NC, 2022.

[13] C. Vela, R. Opromolla, G. Fasano, and H. Schaub, "Modal Solutions-based Approach to Relative Dynamics in the Cislunar Environment," *AIAA SCITECH 2025 Forum*, 2025, p. 2601.

[14] A. W. Koenig and S. D'Amico, "Fast Algorithm for Fuel-Optimal Impulsive Control of Linear Systems With Time-Varying Cost," *IEEE Transactions on Automatic Control*, Vol. 66, No. 9, 2021, pp. 4029–4042, 10.1109/TAC.2020.3027804.

[15] M. Hunter and S. D'Amico, "Fast Algorithm for Fuel-Optimal Constrained Impulsive Control with Application to Satellite Swarming," *2025 IEEE aerospace conference*, Big Sky, MT, IEEE, 2025.

[16] M. Chernick and S. D'Amico, "Closed-form optimal impulsive control of spacecraft formations using reachable set theory," *Journal of Guidance, Control, and Dynamics*, Vol. 44, No. 1, 2021, pp. 25–44.

[17] K. C. Howell, "Three-dimensional periodic halo orbits," *Celestial mechanics*, Vol. 32, 1984, p. 53.

[18] E. M. Zimovan, "Characteristics and design strategies for near rectilinear halo orbits within the earth-moon system," Master's thesis, Purdue University, 2017.