this: POLITECNICO DI TORINO

Master Degree course in Mechatronic Engineering

Master Degree Thesis

Tracking Invariant Manifolds with Nonlinear Model Predictive Control in the Earth–Moon Restricted 3–Body Problem

Supervisors
Prof. Carlo Novara
Prof. Michele Pagone
Prof.ssa Giordana Bucchioni (UniPi)

Candidate
Danilo Campagnolo

Academic Year 2025-2026

"E quindi uscimmo a riveder le stelle."

Dante Alighieri, *Inferno*, Canto XXXIV

Ai miei genitori.

Abstract

The design of efficient guidance and control strategies for spacecraft operating in multi-body gravitational environments represents a challenging problem in astrodynamics. In particular, the Earth–Moon Circular Restricted Three-Body Problem provides a useful framework for describing the nonlinear motion of a spacecraft under the gravitational influence of the Earth and the Moon. These dynamical regions are of increasing interest for space agencies, since libration point orbits can support long–term space infrastructure, such as lunar gateways and future staging platforms for deep–space exploration. Their strategic location also makes them attractive for missions requiring continuous access to the lunar environment and efficient transfers within the cislunar space. Within this model, periodic orbits around libration points and their associated invariant manifolds define natural transport structures that can be exploited for low-energy trajectory design.

This thesis investigates a manifold-based Nonlinear Model Predictive Control strategy for spacecraft guidance toward an Earth–Moon L_2 halo orbit. The proposed approach combines stable invariant manifolds with an artificial-reference NMPC formulation. Instead of forcing the spacecraft to directly track the target periodic orbit, the controller uses stable manifold trajectories as natural reference candidates. An artificial trajectory is optimized inside the NMPC problem and acts as an intermediate path between the real spacecraft motion and the selected natural reference.

The guidance architecture includes online reference selection, reference advancement along the active manifold trajectory, switching among candidate manifolds, and a coasting logic that allows the spacecraft to exploit the uncontrolled CR3BP dynamics when the motion remains sufficiently close to the selected natural path. The nonlinear dynamics are discretized using a fourth-order Runge–Kutta scheme, and the resulting constrained optimization problem is implemented in CasADi.

Numerical simulations in the Earth–Moon CR3BP show that the proposed strategy is able to guide the spacecraft from a perturbed initial condition toward the target halo orbit while limiting continuous actuation. The control effort is mainly concentrated in localized correction phases, whereas extended portions of the trajectory are characterized by nearly zero input. A comparison with a conventional NMPC formulation based on direct halo-orbit tracking highlights a significant reduction in the accumulated Δv, confirming the benefit of exploiting the stable manifold structure within the predictive control framework.

Contents

List of Figures

List of Tables

List of Acronyms

CR3BP	Circular Restricted Three-Body Problem
RTBP	Restricted Three-Body Problem
EM	Earth–Moon
EML2	Earth–Moon L_2
NMPC	Nonlinear Model Predictive Control
MPC	Model Predictive Control
STM	State Transition Matrix
LVLH	Local Vertical Local Horizontal
NLP	Nonlinear Programming Problem
DARE	Discrete Algebraic Riccati Equation
RK4	fourth-order Runge–Kutta
ODE	Ordinary Differential Equation
DU	Distance Unit
TU	Time Unit

Chapter 1

Introduction

In recent decades, interest in cislunar and deep-space missions has grown significantly, both for scientific applications and for human exploration architectures. In this context, the Lagrange points of the EM system represent dynamically strategic regions, capable of offering unique energy and geometric advantages.

In particular, the EML2 point of the EM system is a location of great interest for future space infrastructure, astronomical observation platforms, and logistics hubs for lunar and interplanetary missions. Its location beyond the Moon allows natural access to interplanetary trajectories and a favorable configuration for long-duration cislunar missions. However, orbits around the collinear points of the three-body problem have a fundamental characteristic: intrinsic dynamical instability.

Periodic orbits around EML2, such as Halo or Lissajous orbits, are associated with hyperbolic structures in phase space. This property implies that small perturbations-due to navigation errors, model uncertainties, or environmental disturbances-can grow exponentially over time. As a result, orbital maintenance (station-keeping) and precise tracking of reference trajectories require the use of active control strategies.

At the same time, the nonlinear dynamics of the CR3BP reveal the existence of stable and unstable invariant manifolds associated with periodic orbits. These geometric structures, often interpreted as *dynamical tubes* in phase space, allow low-energy transfers and are a natural resource for trajectory design.

Their presence is not only an element of complexity but also a design opportunity, as it allows the intrinsic dynamics of the system to be exploited to reduce propulsion costs and improve maneuver efficiency.

The scientific interest of this work therefore stems from the intersection of three fundamental aspects: the unstable nature of orbits around EML2, the geometric structure of the invariant manifolds of the CR3BP, and the need to develop advanced control strategies capable of operating in a highly nonlinear and constrained context.

The problem of tracking an unstable periodic orbit in the EM system presents a series of critical issues that cannot be overlooked in a rigorous analysis. The local instability of the orbit implies the presence of directions in phase space along which the error grows exponentially. A direct tracking approach, aimed at rigidly minimizing the distance between the vehicle state and the nominal trajectory, therefore tends to generate high control demands, especially in the presence of perturbations or model uncertainties.

Linear control techniques based on local linearizations are often insufficient in this scenario, as the dynamics of the CR3BP are highly nonlinear and the region of validity of linear approximations can be limited. Furthermore, the presence of constraints on the actuators and the need to reduce propellant consumption require an approach that naturally integrates constraint management and optimization over finite time horizons.

Nonlinear Model Predictive Control is particularly well suited for this purpose, as it allows the incorporation of the complete dynamical model, the management of constraints on states and controls, and the formulation of an optimization problem that takes multiple objectives into account. However, a standard NMPC formulation based on rigid tracking of an unstable trajectory can lead to feasibility problems and a significant increase in propellant consumption, especially when the system deviates from the nominal trajectory along dynamically critical directions.

The central idea behind this work is to reinterpret orbit instability not exclusively as a phenomenon to be counteracted, but as a structural feature to be integrated into the control strategy. Halo orbits around EML2 are associated with stable manifolds that define natural trajectories of approach to the orbit itself. Moving along these directions leads to a spontaneous evolution towards the orbital region, reducing the need for active intervention.

From this perspective, control does not necessarily have to impose rigid tracking of the nominal trajectory, but can be formulated so as to allow flexibility in the definition of the reference. The introduction of an artificial reference within the optimization problem allows the reference itself to be treated as a decision variable, reducing the rigidity of direct tracking while still penalizing deviations from the desired natural structure.

This approach allows the controller to dynamically adapt the reference trajectory according to the current state, exploiting the structure of the dynamical field and reducing the control effort in naturally favorable directions. The integration between the geometry of invariant manifolds and the flexibility offered by NMPC leads to a strategy that coordinates with the dynamics of the system rather than systematically opposing it.

This work aims to develop an integrated methodological framework for tracking periodic orbits around EML2 using a nonlinear predictive control approach enriched by the introduction of an artificial reference. To this end, the dynamics of the EM system are first modeled in the context of the CR3BP, with particular attention to the properties of periodic orbits and their invariant manifolds. Subsequently, a nonlinear control problem is formulated that takes into account dynamic and actuation constraints, as well as the objectives of minimizing propellant consumption and tracking error.

The proposed approach is evaluated through numerical simulations and compared with a conventional NMPC scheme. The assessment considers the resulting closed-loop trajectories, tracking accuracy, control effort, accumulated Δv, coasting phases, switching behavior, and sensitivity to variations in the initial conditions.

1.1 Thesis Organization

The thesis is organized in a progressive way, starting from the dynamical model of the Earth–Moon system and gradually moving toward the development, integration, and validation of the proposed control strategy.

Chapter 2 introduces the dynamical model used throughout the thesis. The CR3BP is presented together with its main assumptions, the synodic reference frame, the non-dimensional formulation, the equations of motion, and the first-order state-space representation. The chapter also discusses the Jacobi integral, the libration points, the controlled dynamics, the discrete-time propagation scheme, and the role of the dynamical model in the subsequent control formulation.

Chapter 3 is devoted to invariant manifolds and dynamical structures in the CR3BP. The chapter first introduces halo orbit families around the Earth–Moon L_2 point and discusses their stability properties. It then presents the local dynamics, the State Transition Matrix, periodic orbits, and the monodromy matrix. Stable and unstable invariant manifolds are defined and interpreted geometrically, with particular emphasis on their role as natural transport structures. Finally, the numerical construction of manifolds and the generation of the stable manifold library are described.

Chapter 4 introduces the Nonlinear Model Predictive Control framework. The chapter presents the general principle of MPC and its nonlinear extension, followed by the finite-horizon optimal control formulation, the cost function, the constraints, and the receding-horizon implementation. The chapter also connects the NMPC formulation to the CR3BP guidance problem and describes the numerical implementation using CasADi.

Chapter 5 presents the artificial reference and structure-aware NMPC formulation. The artificial reference is introduced to provide additional flexibility with respect to rigid trajectory tracking and to create a bridge between the controlled spacecraft trajectory and the natural reference structures.

Chapter 6 describes the integration of invariant manifolds and NMPC. In this chapter, the stable manifold library is incorporated into the predictive control framework, providing dynamically consistent candidate references for the controller.

Chapter 7 presents the numerical setup and the results of the closed-loop simulations. The chapter first introduces the simulation parameters, controller settings, initial conditions, selected manifold references, and relevant implementation choices. The performance of the proposed manifold-based artificial-reference NMPC strategy is then evaluated through the analysis of the spacecraft trajectories, tracking errors, control profiles, accumulated Δv, manifold-switching events, coasting phases, Jacobi constant evolution, and LVLH relative-motion representations.

The proposed strategy is further assessed against a conventional NMPC controller that tracks the target periodic orbit directly, without relying on stable manifold references or an artificial trajectory. The comparison is intended to evaluate the contribution of these additional elements to the resulting closed-loop behavior.

Finally, Chapter 8 summarizes the main conclusions of the thesis and outlines possible future developments.

Chapter 2

Dynamical Model

2.1 Circular Restricted Three-Body Problem

The dynamical environment considered in this work is modeled through the CR3BP. This model provides a suitable framework for describing the motion of a spacecraft under the gravitational attraction of two dominant celestial bodies, called primaries. In the present thesis, the primaries are the Earth and the Moon, while the spacecraft is assumed to have negligible mass with respect to both of them.

Let the masses of the two primaries be denoted by m_1 and m_2, with $m_1 > m_2$. In the CR3BP, the two primaries are assumed to move on circular orbits around their common barycenter. The third body, namely the spacecraft, does not influence the motion of the primaries. This assumption is well justified for space mission applications, since the spacecraft mass is many orders of magnitude smaller than the masses of the Earth and the Moon.

Although simplified with respect to a full ephemeris model, the CR3BP captures the main nonlinear features of the Earth–Moon gravitational environment. In particular, it describes the existence of libration points, periodic orbits, invariant manifolds, and low-energy transfer structures, which are central concepts in dynamical-systems approaches to space mission design [2]. These properties make it especially useful for preliminary mission analysis and for the design of advanced guidance and control strategies.

2.2 Synodic Reference Frame

The model is formulated in a rotating reference frame, also known as the synodic frame (fig. 2.1). This frame rotates with the same angular velocity as the primaries around their barycenter. As a result, the Earth and the Moon remain fixed on the x-axis of the rotating frame.

The origin of the synodic frame is placed at the barycenter of the system. The x-axis is aligned with the line connecting the two primaries, the y-axis lies in the orbital plane, and the z-axis is normal to that plane.

The use of the synodic frame is particularly convenient because it transforms the motion of the primaries into a fixed configuration. Consequently, the equations of motion become autonomous, and equilibrium points can be represented as fixed points in the rotating frame.

In this frame, the gravitational attraction of the primaries is complemented by apparent inertial effects due to rotation, namely the Coriolis and centrifugal accelerations.

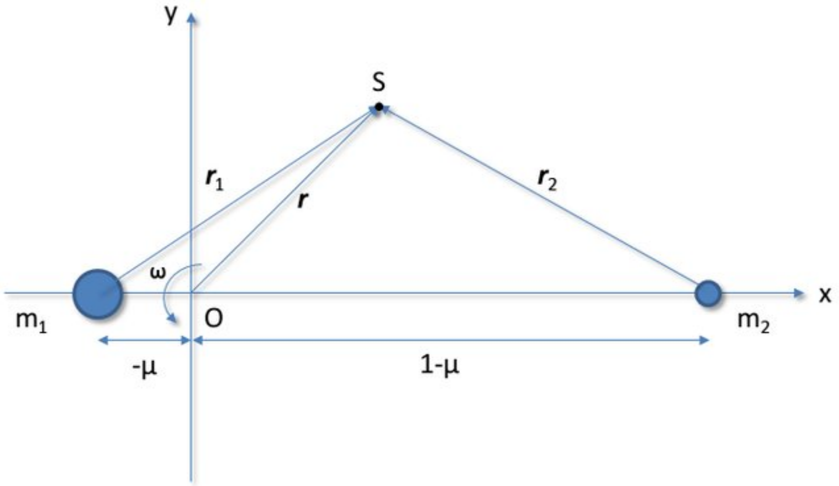

Figure 2.1: Synodic reference frame adopted in the CR3BP. The origin is located at the barycenter of the system, while the two primaries remain fixed along the x-axis. [Reproduced from Sullo et al. [1]]

2.3 Non-dimensional Formulation

A non-dimensional formulation is adopted in order to obtain compact equations of motion and to reduce the number of physical parameters appearing in the model. The characteristic distance is chosen as the mean distance between the Earth and the Moon, while the characteristic time is derived from the sidereal period of the Moon around the Earth–Moon barycenter.

The Moon orbits the Earth–Moon barycenter with a sidereal period of approximately 27.322 days, corresponding to 27 days, 7 hours, 43 minutes, and 11.5 seconds. This dimensional period is denoted by P_{EM}. The distance unit is defined as

$$DU = d_{EM}, \tag{2.1}$$

where d_{EM} is the mean Earth–Moon distance. In this work, the value

$$DU = 384400 \times 10^3 \text{ m} \tag{2.2}$$

is adopted.

The time unit is chosen so that the non-dimensional angular velocity of the rotating Earth–Moon system is equal to unity. Since one complete revolution corresponds to an angle of 2π radians, the time unit is defined as

$$TU = \frac{P_{EM}}{2\pi}. \tag{2.3}$$

Therefore, one non-dimensional time unit corresponds to the time required by the Earth–Moon rotating frame to sweep an angle of one radian. With this choice, a complete lunar revolution corresponds to a non-dimensional time interval of 2π.

In this work, the sidereal period is taken as

$$P_{EM} = 27.321661 \cdot 24 \cdot 60 \cdot 60 \text{ s}, \tag{2.4}$$

and therefore

$$TU = \frac{27.321661 \cdot 24 \cdot 60 \cdot 60}{2\pi} \text{ s}. \tag{2.5}$$

The dimensional and non-dimensional variables are related by

$$\mathbf{r} = \frac{\mathbf{r}_{dim}}{DU}, \qquad t = \frac{t_{dim}}{TU}, \qquad \mathbf{v} = \frac{\mathbf{v}_{dim}}{DU/TU}, \qquad \mathbf{a} = \frac{\mathbf{a}_{dim}}{DU/TU^2}. \tag{2.6}$$

With this normalization, the Earth–Moon system is completely characterized by the mass parameter

$$\mu = \frac{m_2}{m_1 + m_2}. \tag{2.7}$$

For the Earth–Moon system, m_1 represents the Earth and m_2 represents the Moon. The non-dimensional parameter μ determines the relative location of the two primaries in the synodic frame. In particular, the larger primary is located at

$$\mathbf{r}_1 = [-\mu, 0, 0]^T, \tag{2.8}$$

whereas the smaller primary is located at

$$\mathbf{r}_2 = [1 - \mu, 0, 0]^T. \tag{2.9}$$

Therefore, in non-dimensional coordinates, the distance between the two primaries is equal to one. This normalization is useful not only from an analytical point of view, but also from a numerical perspective, since it avoids the simultaneous presence of very large distances and very small accelerations in the equations of motion.

2.4 Equations of Motion

The spacecraft state vector is defined as

$$\mathbf{x} = \begin{bmatrix} x & y & z & \dot{x} & \dot{y} & \dot{z} \end{bmatrix}^T, \tag{2.10}$$

where (x, y, z) are the non-dimensional position coordinates in the synodic frame and $(\dot{x}, \dot{y}, \dot{z})$ are the corresponding velocity components.

The distances of the spacecraft from the two primaries are given by

$$r_1 = \sqrt{(x + \mu)^2 + y^2 + z^2}, \tag{2.11}$$

and

$$r_2 = \sqrt{(x - 1 + \mu)^2 + y^2 + z^2}. \tag{2.12}$$

The effective potential of CR3BP is defined as the sum:

$$U^* = U + U_C, \tag{2.13}$$

where U is the gravitational term, whereas U_C is the centrifugal term. So, substituting the various terms yields:

$$U^*(x, y, z) = \frac{1-\mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}\left(x^2 + y^2\right). \tag{2.14}$$

Using this potential, the uncontrolled equations of motion can be written as

$$\ddot{x} = 2\dot{y} + \frac{\partial U^*}{\partial x}, \tag{2.15}$$

$$\ddot{y} = -2\dot{x} + \frac{\partial U^*}{\partial y}, \tag{2.16}$$

$$\ddot{z} = \frac{\partial U^*}{\partial z}. \tag{2.17}$$

Expanding the partial derivatives, the equations become

$$\ddot{x} = 2\dot{y} + x - \frac{1-\mu}{r_1^3}(x+\mu) - \frac{\mu}{r_2^3}(x-1+\mu), \tag{2.18}$$

$$\ddot{y} = -2\dot{x} + y - \frac{1-\mu}{r_1^3}y - \frac{\mu}{r_2^3}y, \tag{2.19}$$

$$\ddot{z} = -\frac{1-\mu}{r_1^3}z - \frac{\mu}{r_2^3}z. \tag{2.20}$$

The terms $2\dot{y}$ and $-2\dot{x}$ are the Coriolis contributions, while the terms depending directly on x and y include the centrifugal acceleration. The remaining terms represent the gravitational attraction of the two primaries.

These equations are nonlinear and strongly coupled. This feature is central to the present work, since the same nonlinear dynamics that make the control problem challenging also generate the invariant manifold structures exploited for low-energy transfers.

2.5 First-Order State-Space Form

For numerical propagation and later control design, the second-order equations of motion are rewritten in first-order state-space form as

$$\dot{\mathbf{x}} = f(\mathbf{x}). \tag{2.21}$$

More explicitly, the natural CR3BP dynamics are given by

$$\dot{\mathbf{x}} = \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \\ 2\dot{y} + x - \frac{1-\mu}{r_1^3}(x+\mu) - \frac{\mu}{r_2^3}(x-1+\mu) \\ -2\dot{x} + y - \frac{1-\mu}{r_1^3}y - \frac{\mu}{r_2^3}y \\ -\frac{1-\mu}{r_1^3}z - \frac{\mu}{r_2^3}z \end{bmatrix}. \tag{2.22}$$

This compact representation is the baseline dynamical model used throughout the thesis. In the following sections, its main dynamical properties are first discussed before introducing the controlled form of the equations.

2.6 Jacobi Integral

For the uncontrolled CR3BP, the system admits a conserved quantity known as the Jacobi constant. It is defined as

$$C = 2U^*(x, y, z) - \left(\dot{x}^2 + \dot{y}^2 + \dot{z}^2\right). \tag{2.23}$$

The Jacobi constant can be interpreted as an energy-like quantity in the rotating frame. Since it is conserved along natural trajectories, it provides important information about the regions of space that are dynamically accessible to the spacecraft.

In particular, from the definition of C one obtains

$$\dot{x}^2 + \dot{y}^2 + \dot{z}^2 = 2U^*(x, y, z) - C. \tag{2.24}$$

Since the squared velocity must be non-negative, the motion is possible only in the regions satisfying

$$2U^*(x, y, z) - C \geq 0. \tag{2.25}$$

The equality

$$2U^*(x, y, z) - C = 0 \tag{2.26}$$

defines the so-called zero-velocity surfaces, which separate the admissible regions of motion, also known as Hill's regions, from the forbidden regions for a given value of the Jacobi constant.

In the context of this thesis, the Jacobi constant is relevant because it characterizes natural trajectories, periodic orbits, and their associated invariant manifolds. Moreover, it will later be used as a diagnostic quantity in closed-loop simulations to evaluate how the controlled trajectory departs from or remains close to the selected natural reference.

2.7 Libration Points

The CR3BP admits five equilibrium points, commonly referred to as *Lagrange points* or *libration points*. These points correspond to fixed positions in the synodic reference frame, meaning that a particle placed exactly at one of them with zero velocity remains at rest with respect to the rotating frame.

The libration points are obtained by imposing zero velocity and zero acceleration in the equations of motion. Since the dynamics can be expressed in terms of the pseudo-potential function U^*, the equilibrium condition is written as

$$\nabla U^* = 0. \tag{2.27}$$

In component form, this condition corresponds to

$$\frac{\partial U^*}{\partial x} = 0, \qquad \frac{\partial U^*}{\partial y} = 0, \qquad \frac{\partial U^*}{\partial z} = 0. \tag{2.28}$$

All five libration points lie in the plane of motion of the primaries, and therefore satisfy $z = 0$. Three of them, denoted by L_1, L_2, and L_3, are located on the x-axis and are *called collinear libration points*. The remaining two, denoted by L_4 and L_5, form equilateral triangles with the two primaries and are called *triangular libration points*.

For the collinear libration points, the equilibrium condition reduces to a scalar nonlinear equation along the x-axis:

$$x - \frac{(1-\mu)(x+\mu)}{|x+\mu|^3} - \frac{\mu(x-1+\mu)}{|x-1+\mu|^3} = 0, \tag{2.29}$$

with $y = z = 0$.

The three real solutions of Eq. (2.29) correspond to the collinear libration points L_1, L_2, and L_3. Since this equation is nonlinear, the locations of the collinear points are computed numerically. In this work, MATLAB's `fzero` function is used to solve Eq. (2.29) in three different intervals, each containing one of the collinear points:

$$L_1 \in (0{,}1-\mu), \qquad L_2 \in (1-\mu, +\infty), \qquad L_3 \in (-\infty, -\mu). \tag{2.30}$$

More specifically, the numerical search intervals are selected as

$$L_1 \in [0.01{,}0.97], \qquad L_2 \in [1{,}1.5], \qquad L_3 \in [-1.5, -1]. \tag{2.31}$$

These intervals isolate the three roots of the collinear equilibrium equation for the Earth–Moon system. In the adopted non-dimensional synodic frame, L_1 is located between the Earth and the Moon, L_2 lies beyond the Moon, and L_3 lies on the opposite side of the Earth with respect to the Moon.

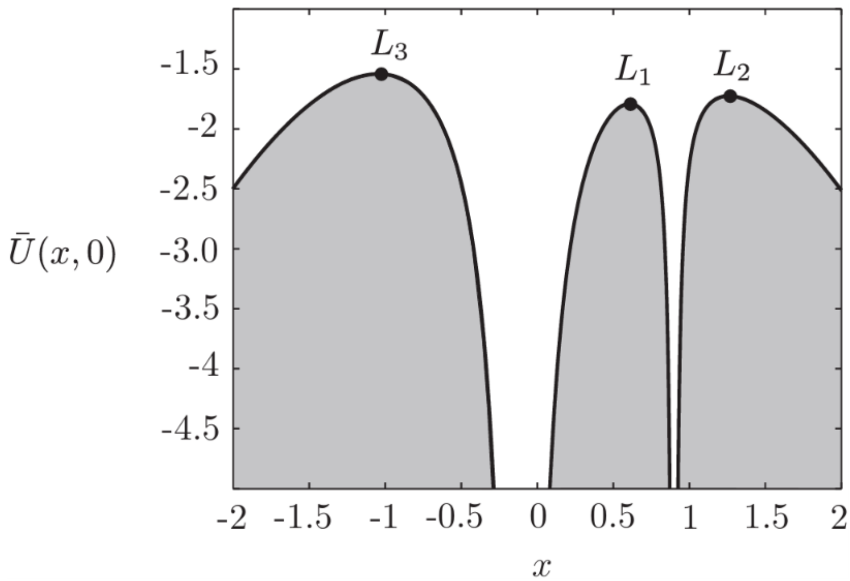

Figure 2.2: Profile of $U^*(x{,}0)$ for $\mu = 0.1$, obtained as the intersection of the pseudo-potential surface with the plane $y = 0$. At the x-locations of the two primaries, m_1 and m_2, the function tends to $-\infty$, following the sign convention adopted in Koon et al. [2], while the local maxima correspond to the unstable collinear libration points L_1, L_2, and L_3. [Reproduced from Koon et al. [2]]

The triangular points, instead, have an explicit geometrical expression. In non-dimensional coordinates, their positions are

$$L_4 = \left[\frac{1}{2} - \mu,\, \frac{\sqrt{3}}{2},\, 0\right]^T, \qquad L_5 = \left[\frac{1}{2} - \mu,\, -\frac{\sqrt{3}}{2},\, 0\right]^T. \tag{2.32}$$

These points are symmetric with respect to the x-axis and are located at the vertices of equilateral triangles formed with the two primaries.

The stability properties of the libration points depend on the linearized dynamics around each equilibrium. The collinear points L_1, L_2, and L_3 are unstable in the CR3BP. The triangular points L_4 and L_5 may be stable or unstable depending on the value of the mass parameter μ. In the Earth–Moon system, they satisfy the classical linear stability condition.

Among the collinear points, L_2 is the one considered in this work. Its associated periodic orbits and invariant manifolds provide the dynamical structures used in the following chapters for trajectory generation and control design.

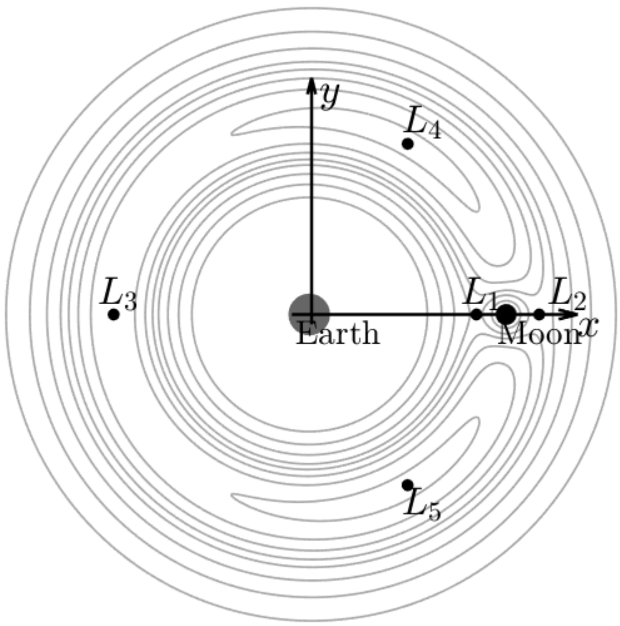

Figure 2.3: Earth–Moon CR3BP in the synodic frame, showing the five libration points and the Jacobi constant contours associated with the zero-velocity curves.

2.8 Controlled Dynamics

After defining the natural CR3BP model and its main dynamical properties, the effect of spacecraft actuation can be introduced. In the control problem, the spacecraft is assumed to be equipped with a low-thrust propulsion system. The control input is modeled as an acceleration vector

$$\mathbf{u} = \begin{bmatrix} u_x & u_y & u_z \end{bmatrix}^T. \tag{2.33}$$

The controlled equations of motion are therefore obtained by adding the control components directly to the acceleration equations:

$$\ddot{x} = 2\dot{y} + x - \frac{1-\mu}{r_1^3}(x+\mu) - \frac{\mu}{r_2^3}(x-1+\mu) + u_x, \tag{2.34}$$

$$\ddot{y} = -2\dot{x} + y - \frac{1-\mu}{r_1^3}y - \frac{\mu}{r_2^3}y + u_y, \tag{2.35}$$

$$\ddot{z} = -\frac{1-\mu}{r_1^3}z - \frac{\mu}{r_2^3}z + u_z. \tag{2.36}$$

In compact first-order form, the controlled dynamics can be written as

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}), \tag{2.37}$$

where

$$f(\mathbf{x}, \mathbf{u}) = \begin{bmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \\ 2\dot{y} + U_x^* + u_x \\ -2\dot{x} + U_y^* + u_y \\ U_z^* + u_z \end{bmatrix}. \tag{2.38}$$

This formulation is consistent with continuous low-thrust control, where the input acts as an acceleration applied to the spacecraft. In the absence of control, namely for $\mathbf{u} = \mathbf{0}$, the system reduces to the natural CR3BP dynamics previously introduced.

When control is applied, the Jacobi constant is no longer conserved. This is expected, since the thrust acceleration modifies the mechanical energy of the spacecraft. Nevertheless, the Jacobi constant remains useful as a diagnostic quantity for understanding the dynamical behavior of the controlled trajectory.

2.9 Discrete-Time Propagation

The continuous-time equations of motion must be discretized for numerical simulations and for the NMPC formulation introduced in the following chapters. In this work, a fourth-order Runge–Kutta scheme is adopted due to its favorable compromise between numerical accuracy and computational efficiency.

Given the controlled continuous-time system

$$\dot{\mathbf{x}} = f(\mathbf{x}, \mathbf{u}), \tag{2.39}$$

and a sampling time Δt, the RK4 method computes the following intermediate slopes:

$$k_1 = f(\mathbf{x}_k, \mathbf{u}_k), \tag{2.40}$$

$$k_2 = f\left(\mathbf{x}_k + \frac{\Delta t}{2}k_1, \mathbf{u}_k\right), \tag{2.41}$$

$$k_3 = f\left(\mathbf{x}_k + \frac{\Delta t}{2}k_2, \mathbf{u}_k\right), \tag{2.42}$$

$$k_4 = f\left(\mathbf{x}_k + \Delta t k_3, \mathbf{u}_k\right). \tag{2.43}$$

The state update is then given by

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \frac{\Delta t}{6}\left(k_1 + 2k_2 + 2k_3 + k_4\right). \tag{2.44}$$

In the numerical implementation, the propagation over one sampling interval is further subdivided into M internal integration steps. If T_s denotes the sampling time and M is the number of RK4 substeps, the internal integration step is

$$h = \frac{T_s}{M}. \tag{2.45}$$

This discretization is used consistently both for the propagation of trajectories and inside the prediction model of the controller. This consistency is important because it avoids artificial discrepancies between the reference trajectories and the dynamics used by the optimizer.

The RK4 scheme is particularly suitable in the present context because the CR3BP is strongly nonlinear and sensitive to initial conditions, especially in the vicinity of unstable libration point orbits.

2.9.1 Numerical Verification of the RK4 Discretization

To verify the accuracy of the discrete-time propagation scheme, the RK4 integrator is compared with a high-accuracy continuous-time reference solution obtained using the MATLAB adaptive solver `ode113`. The comparison is performed for different values of the number of internal RK4 substeps M.

The propagation accuracy is evaluated by considering the full-state error, the position error, and the velocity error:

$$e_x(k) = \left\|\mathbf{x}_{RK4,k} - \mathbf{x}_{ref,k}\right\|, \tag{2.46}$$
$$e_r(k) = \left\|\mathbf{r}_{RK4,k} - \mathbf{r}_{ref,k}\right\|, \tag{2.47}$$
$$e_v(k) = \left\|\mathbf{v}_{RK4,k} - \mathbf{v}_{ref,k}\right\|. \tag{2.48}$$

Figures 2.4, 2.5, and 2.6 show the evolution of the full-state, position, and velocity propagation errors for different values of M, respectively. As expected, increasing the number of internal RK4 substeps reduces the propagation error.

It can also be observed that the propagation error tends to increase over time. This behavior is expected, since the local truncation errors introduced at each integration step accumulate during the propagation. Moreover, the CR3BP dynamics are sensitive to small perturbations, especially near unstable libration-point trajectories and dynamically demanding regions. As a result, even small numerical discrepancies between the RK4 solution and the high-accuracy reference solution may be amplified as the propagation time increases.

This confirms that the selected discretization provides a consistent numerical approximation of the continuous-time CR3BP dynamics, while also highlighting the importance of using a sufficiently accurate integration scheme in this sensitive dynamical environment.

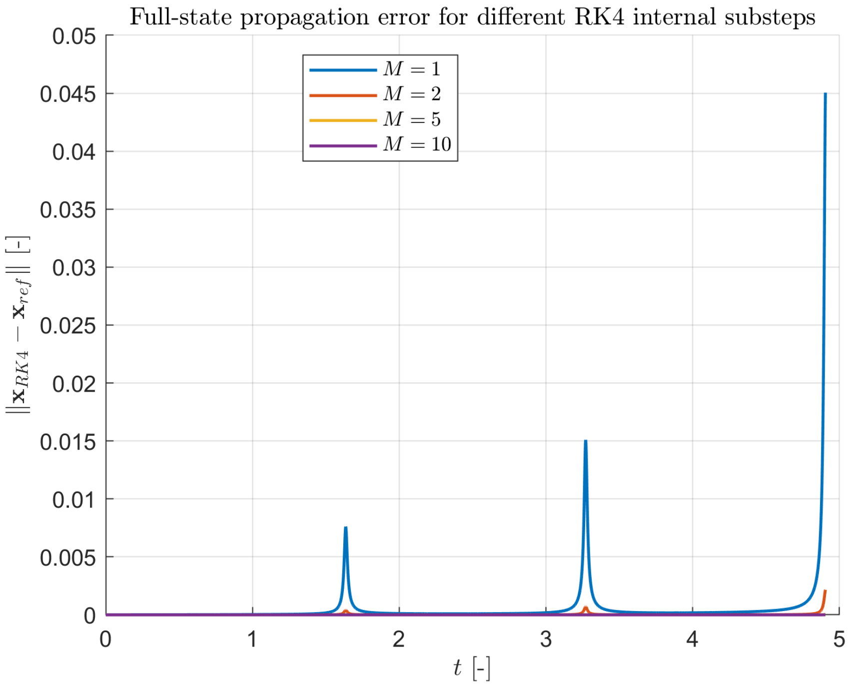

Figure 2.4: Full-state propagation error for different RK4 internal substeps M

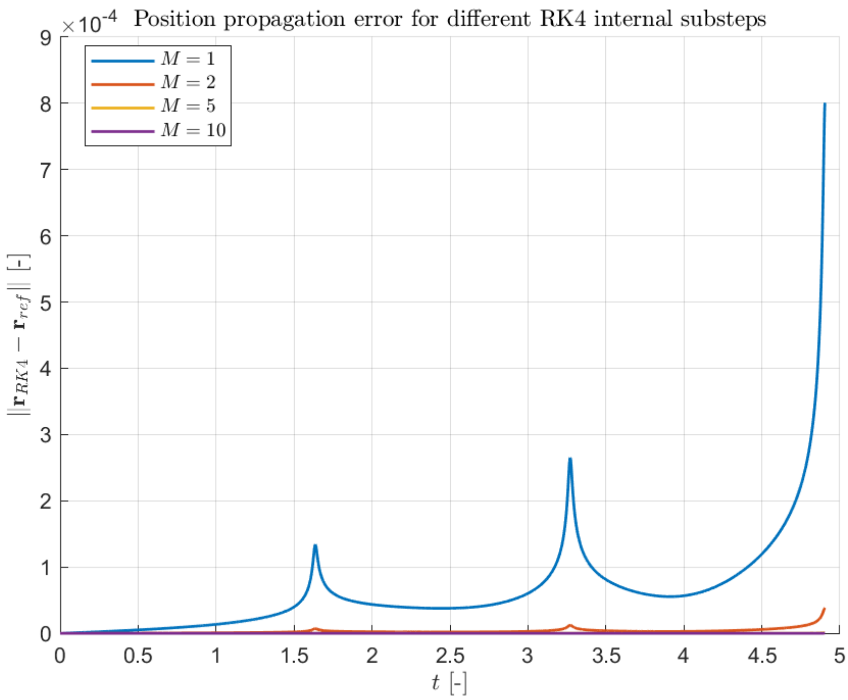

Figure 2.5: Position propagation error for different RK4 internal substeps M

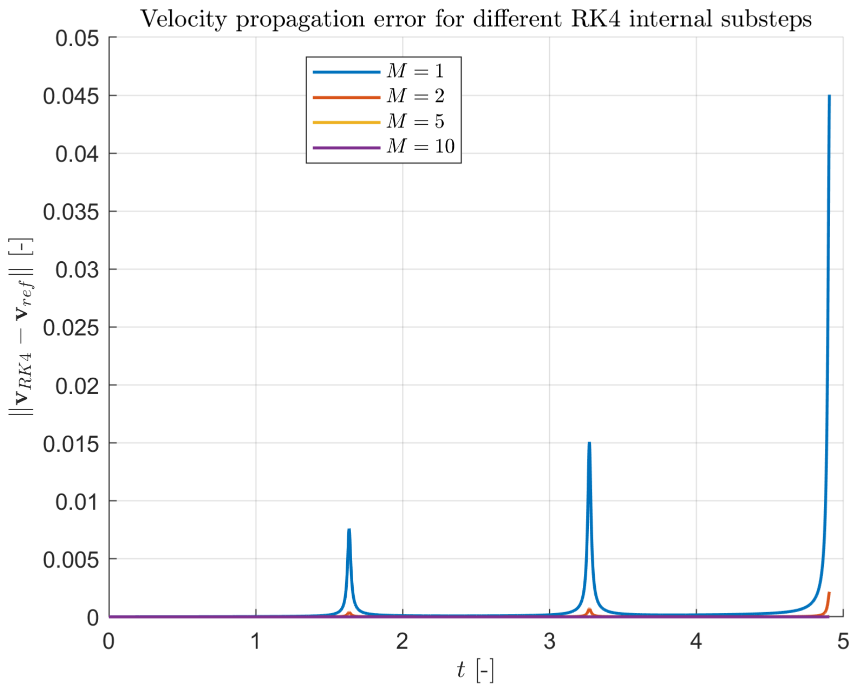

Figure 2.6: Velocity propagation error for different RK4 internal substeps M

2.10 Concluding Remarks

This chapter introduced the dynamical model adopted throughout the thesis. The EM system was described within the framework of the CR3BP, using the synodic reference frame and a non-dimensional formulation. The equations of motion were derived in first-order state-space form, and the Jacobi integral was introduced as a fundamental quantity characterizing the natural dynamics of the system.

The libration points identified by the CR3BP provide the basis for the periodic orbits discussed in the following chapter. In particular, the instability of periodic solutions around the collinear libration points naturally leads to the invariant manifolds that will be introduced next.

Finally, the controlled form of the equations of motion and the discrete-time propagation scheme were presented. These elements provide the connection between the theoretical description of the EM environment and the numerical implementation of the guidance and control strategy developed in the remainder of the thesis.

Chapter 3

Invariant Manifolds and Dynamical Structures

3.1 Introduction

As discussed in the previous chapter, the dynamics of the CR3BP are characterized by the presence of equilibrium points and families of periodic orbits. However, these objects alone do not fully describe the structure of the motion in phase space.

A deeper understanding of the system can be achieved by introducing invariant manifolds, which represent geometric structures that organize the global behavior of trajectories. These manifolds can be interpreted as pathways along which motion naturally evolves, providing a powerful framework for analyzing and exploiting the nonlinear dynamics of the CR3BP.

3.2 Halo Orbit Family and Stability Index

Periodic orbits are recurrent solutions of the CR3BP equations of motion. A trajectory is periodic if, after a finite time interval, it returns to the same state in phase space. If $\mathbf{x}^*(t)$ denotes a periodic solution, this condition can be written as

$$\mathbf{x}^*(t+T) = \mathbf{x}^*(t), \tag{3.1}$$

where T is the period of the orbit.

Around the collinear libration points, the CR3BP admits several families of periodic orbits. Among them, halo orbits are three-dimensional periodic solutions that extend out of the orbital plane of the primaries. These orbits have been extensively studied both analytically and numerically [4, 5], and they do not appear as isolated trajectories, but rather as continuous families whose geometry, amplitude, period, and stability properties vary from one member to another.

The stability of a periodic orbit describes the local behavior of nearby trajectories. In general, an unstable periodic orbit is characterized by directions in phase space along which small perturbations grow with time. To provide a compact measure of this behavior, a stability index ν can be associated with each periodic orbit.

This stability index is defined as

$$\nu = \frac{1}{2}\left(\lambda_u + \frac{1}{\lambda_u}\right), \tag{3.2}$$

where λ_u denotes the unstable multiplier associated with the periodic orbit. The formal definition of this multiplier is introduced in the following sections through the linearized dynamics and the monodromy matrix. Values of ν close to unity indicate weak instability, whereas larger values correspond to stronger sensitivity to perturbations.

For visualization purposes, the logarithmic quantity $\log_{10}(\nu)$ is often more convenient than ν itself, since the stability index may vary over a wide range along the same family. Figure 3.1 illustrates an L_2 halo orbit family in the Earth–Moon CR3BP, with the orbits colored according to $\log_{10}(\nu)$. The figure shows that both the geometry and the stability properties of the periodic solutions change continuously along the family.

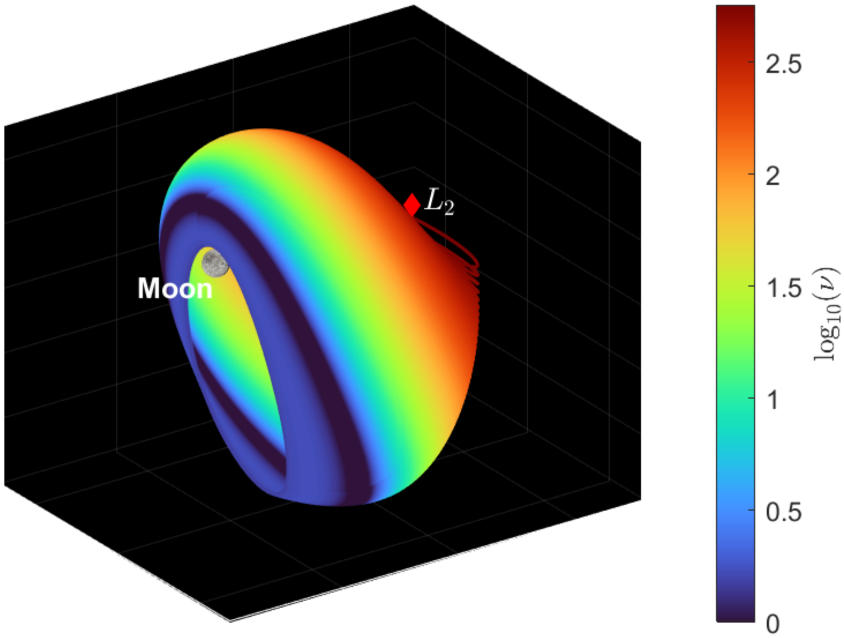

Figure 3.1: L_2 halo orbit family in the Earth–Moon CR3BP, colored according to $\log_{10}(\nu)$, where ν is the stability index associated with each periodic orbit.

This visualization provides a qualitative introduction to the link between periodic orbits and stability. The formal tools used to quantify this behavior, namely the linearized dynamics, the State Transition Matrix, and the monodromy matrix, are introduced in the following sections.

3.3 Local Dynamics and Linearization

In order to analyze the behavior of trajectories in the neighborhood of a reference solution, it is useful to study the local dynamics of small perturbations. Let

$$\mathbf{x}^*(t) \tag{3.3}$$

be a nominal solution of the nonlinear CR3BP equations of motion. In the context of this work, this reference solution is typically a periodic orbit around the Earth–Moon L_2 point.

A neighboring trajectory can be written as

$$\mathbf{x}(t) = \mathbf{x}^*(t) + \delta\mathbf{x}(t), \tag{3.4}$$

where $\delta\mathbf{x}(t)$ represents a small perturbation with respect to the nominal trajectory.

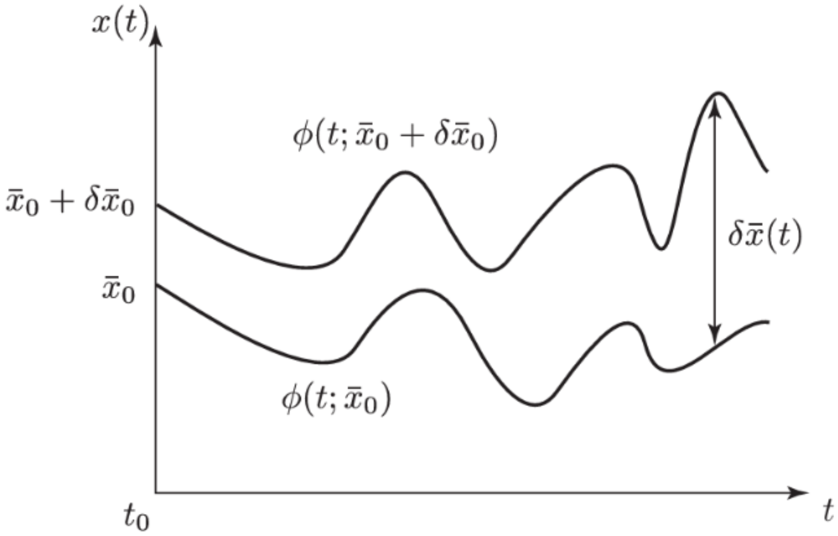

Figure 3.2: A trajectory $\phi(t; \bar{x}0)$ and a neighboring trajectory, $\phi(t; \bar{x}0+\delta\bar{x}0)$. [Reproduced from Koon et al. [2]]

Substituting this expression into the nonlinear dynamics

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) \tag{3.5}$$

and retaining only the first-order terms yields the linearized system

$$\delta\dot{\mathbf{x}}(t) = A(t)\delta\mathbf{x}(t), \tag{3.6}$$

where

$$A(t) = \left.\frac{\partial \mathbf{f}}{\partial \mathbf{x}}\right|_{\mathbf{x}^*(t)} \tag{3.7}$$

is the Jacobian matrix of the CR3BP vector field evaluated along the reference solution.

Equation (3.6) describes how infinitesimal deviations evolve in the vicinity of the nominal trajectory. Although the original CR3BP dynamics are nonlinear, the linearized system provides a local approximation that is sufficient to identify directions of expansion and contraction in phase space. These directions are fundamental for the study of stability and for the construction of invariant manifolds.

Since the reference solution $\mathbf{x}^*(t)$ may vary with time, the matrix $A(t)$ is generally time-dependent. In particular, when $\mathbf{x}^*(t)$ is a periodic orbit of period T, the Jacobian matrix is also periodic:

$$A(t+T) = A(t). \tag{3.8}$$

This property is the basis for the use of Floquet theory in the stability analysis of periodic orbits.

3.4 State Transition Matrix

The solution of the linearized system can be expressed through the State Transition Matrix, usually denoted by

$$\Phi(t, t_0). \tag{3.9}$$

The STM maps an initial perturbation at time t_0 into the corresponding perturbation at time t. Therefore,

$$\delta\mathbf{x}(t) = \Phi(t, t_0)\delta\mathbf{x}(t_0). \tag{3.10}$$

The State Transition Matrix satisfies the matrix differential equation

$$\dot{\Phi}(t, t_0) = A(t)\Phi(t, t_0), \qquad \Phi(t_0, t_0) = I, \tag{3.11}$$

where I is the identity matrix. In the CR3BP, the STM is a 6×6 matrix, since the state vector contains three position components and three velocity components.

The STM plays a central role in the analysis of the local dynamics. Each column of $\Phi(t, t_0)$ represents the response of the system to a unit perturbation in one component of the initial state. As a result, the STM contains information about the sensitivity of the trajectory to initial conditions.

For numerical applications, the STM is usually propagated together with the nonlinear equations of motion. The combined system is composed of the original CR3BP dynamics and the variational equations:

$$\begin{cases} \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}), \\ \dot{\Phi} = A(t)\Phi. \end{cases} \tag{3.12}$$

This simultaneous propagation makes it possible to obtain both the nominal trajectory and the local linear mapping around it.

3.5 Periodic Orbits and Monodromy Matrix

When the reference trajectory is periodic, the State Transition Matrix evaluated over one complete period provides direct information about the local stability of the orbit. Let

$\mathbf{x}^*(t)$ be a periodic orbit with period T, such that

$$\mathbf{x}^*(t+T) = \mathbf{x}^*(t). \tag{3.13}$$

The STM over one period is defined as

$$M = \Phi(T,0), \tag{3.14}$$

and is called the monodromy matrix. This matrix maps an initial perturbation at a point on the periodic orbit into the corresponding perturbation after one complete revolution:

$$\delta\mathbf{x}(T) = M\delta\mathbf{x}(0). \tag{3.15}$$

The eigenvalues of M are known as Floquet multipliers. They describe the evolution of infinitesimal perturbations from one period to the next. This interpretation follows from Floquet theory, which states that the stability of a periodic linear time-varying system is determined by the eigenvalues of the monodromy matrix. If a multiplier has modulus greater than one, perturbations along the associated eigendirection grow after one revolution. If its modulus is smaller than one, perturbations decay. If its modulus is equal to one, the corresponding direction is neutrally stable.

For halo orbits around collinear libration points, the Floquet multipliers can be grouped into pairs with distinct geometrical meanings. Following the description given by Gómez et al. [6], the typical structure for small and medium-size halo orbits consists of one real hyperbolic pair, one unit pair, and one complex conjugate pair with unit modulus.

The first pair is composed of two real multipliers,

$$\lambda_u > 1, \qquad \lambda_s < 1, \qquad \lambda_u\lambda_s = 1. \tag{3.16}$$

These multipliers are associated with the hyperbolic character of the periodic orbit. The multiplier λ_u defines the most expansive direction, while λ_s defines the contracting direction. The corresponding eigenvectors are denoted here as $\mathbf{v}_u(0)$ and $\mathbf{v}_s(0)$ and represent, respectively, the local unstable and stable directions at the initial point of the orbit.

The second pair is composed of two unit multipliers,

$$\lambda_3 = \lambda_4 = 1. \tag{3.17}$$

One of these directions is tangent to the periodic orbit and is associated with a phase shift along the orbit itself. The other is related to variations along the family of halo orbits, for example variations of the amplitude parameter used to identify different members of the family.

The third pair is composed of two complex conjugate multipliers with unit modulus,

$$\lambda_5 = \overline{\lambda}_6, \qquad |\lambda_5| = |\lambda_6| = 1. \tag{3.18}$$

These multipliers are associated with the center dynamics of the periodic orbit and correspond to bounded oscillatory perturbations in the linear approximation.

The local stability of the halo orbit can therefore be interpreted as the combination of hyperbolic, neutral, and center behavior. The hyperbolic pair is the most relevant for the construction of invariant manifolds, since it defines the local directions along which trajectories approach or depart from the periodic orbit.

The stability index qualitatively introduced in Section 3.2 can now be expressed in terms of the unstable Floquet multiplier as

$$\nu = \frac{1}{2}\left(\lambda_u + \frac{1}{\lambda_u}\right). \tag{3.19}$$

Values of ν close to unity correspond to weakly unstable orbits, whereas larger values indicate stronger hyperbolic behavior and greater sensitivity to perturbations.

The stable and unstable eigendirections are not limited to the initial point of the orbit. They can be transported along the periodic trajectory through the variational flow:

$$\mathbf{v}_u(t) = \Phi(t,0)\mathbf{v}_u(0), \qquad \mathbf{v}_s(t) = \Phi(t,0)\mathbf{v}_s(0). \tag{3.20}$$

In this way, a stable and an unstable direction can be associated with each point of the periodic orbit.

These transported eigendirections provide the local linear approximation of the invariant manifolds. A perturbation along the unstable direction generates trajectories that depart from the periodic orbit in forward time. Conversely, a perturbation along the stable direction generates trajectories that approach the periodic orbit in forward time, or equivalently depart from it when propagated backward in time.

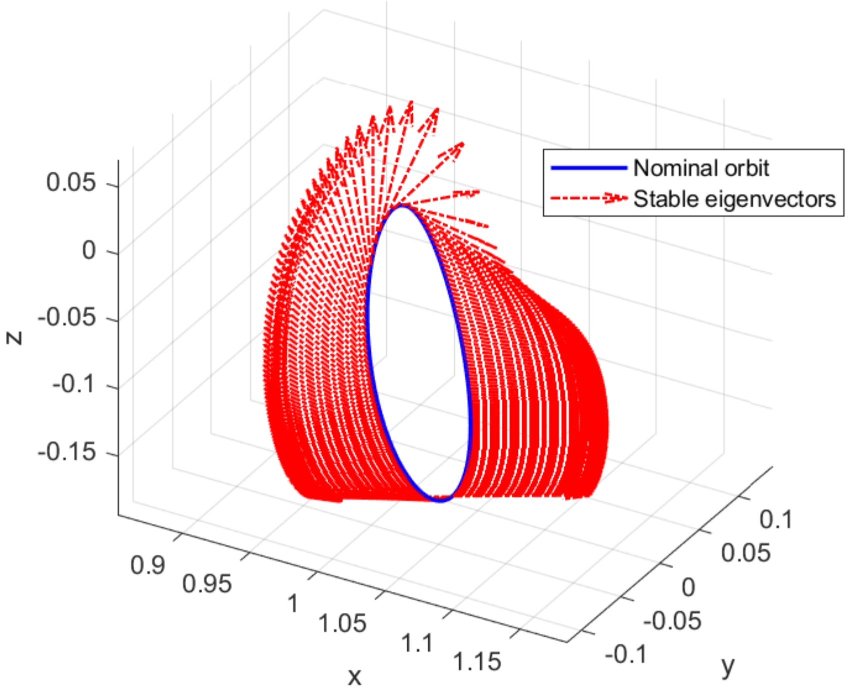

Figure 3.3: Local stable eigendirections computed from the eigenstructure of the monodromy matrix along the periodic orbit.

3.6 Invariant Manifolds

Invariant manifolds are geometric structures associated with equilibrium points or periodic orbits of a dynamical system. In the context of the CR3BP, they play a fundamental role in describing how trajectories evolve in the neighborhood of unstable periodic solutions, such as halo orbits around the collinear libration points.

Let $\mathbf{x}^*(t)$ be a periodic orbit of period T. The stable manifold associated with this orbit is defined as the set of initial conditions whose trajectories asymptotically approach the periodic orbit as time increases:

$$W^s(\mathbf{x}^*) = \left\{ \mathbf{x}_0 \;\middle|\; \lim_{t \to +\infty} \operatorname{dist}\left(\mathbf{x}(t; \mathbf{x}_0), \mathbf{x}^*(t)\right) = 0 \right\}. \tag{3.21}$$

Similarly, the unstable manifold is defined as the set of initial conditions whose trajectories asymptotically approach the periodic orbit backward in time:

$$W^u(\mathbf{x}^*) = \left\{ \mathbf{x}_0 \;\middle|\; \lim_{t \to -\infty} \operatorname{dist}\left(\mathbf{x}(t; \mathbf{x}_0), \mathbf{x}^*(t)\right) = 0 \right\}. \tag{3.22}$$

Equivalently, trajectories belonging to the stable manifold approach the periodic orbit in forward time, whereas trajectories belonging to the unstable manifold depart from it in forward time. The term invariant means that these sets are preserved by the flow of the system: if a trajectory starts on a manifold, it remains on that manifold under the natural evolution of the dynamics.

3.7 Geometrical Interpretation and Role

From a geometrical point of view, invariant manifolds can be interpreted as surfaces embedded in the phase space of the system. Since the CR3BP state is six-dimensional, the complete manifold structure cannot be directly visualized. However, when projected onto configuration space or selected phase-space coordinates, stable and unstable manifolds often appear as tube-like structures attached to the periodic orbit.

These tubes organize the motion in the neighborhood of the orbit. The stable manifold identifies natural directions of approach, while the unstable manifold identifies natural directions of departure. Therefore, invariant manifolds provide a link between the local stability properties of a periodic orbit and the global behavior of trajectories in the nonlinear system.

This interpretation is particularly important in the CR3BP. Although the stable and unstable directions are first obtained from a local linear analysis, the corresponding manifolds become nonlinear structures when propagated under the full equations of motion. As a result, they may extend far from the periodic orbit, bend, fold, and connect different regions of phase space.

For this reason, invariant manifolds can be regarded as natural transport structures. This interpretation is widely used in dynamical-systems approaches to space mission design, where stable and unstable manifolds are exploited to identify natural pathways and low-energy connections in multi-body gravitational environments [2]. A trajectory lying exactly on a manifold evolves according to the uncontrolled CR3BP dynamics. In ideal conditions, no continuous control input is required to follow such a trajectory. This makes stable manifolds especially useful as dynamically consistent reference paths: rather than imposing an arbitrary trajectory, one can exploit a path that is already compatible with the natural dynamics of the system.

In the proposed guidance framework, this property is used by considering stable manifold trajectories as natural reference candidates. The role of the controller is then not to force the spacecraft along an artificial path, but to connect the controlled motion to a dynamically meaningful structure.

3.8 Numerical Construction of Manifolds

The numerical construction of invariant manifolds combines the local stability information obtained from the monodromy matrix with nonlinear propagation under the CR3BP equations of motion.

The first step consists of computing the periodic orbit and its associated stable and unstable eigendirections. Let $\mathbf{v}_s(t_i)$ and $\mathbf{v}_u(t_i)$ denote, respectively, the stable and unstable directions associated with a sampled point $\mathbf{x}^*(t_i)$ on the periodic orbit. A small perturbation is then applied along these directions.

For the stable manifold, the perturbed initial conditions are written as

$$\mathbf{x}_{0,i}^{s,\pm} = \mathbf{x}^*(t_i) \pm \varepsilon \frac{\mathbf{v}_s(t_i)}{\|\mathbf{v}_s(t_i)\|}, \tag{3.23}$$

where ε is a small scalar perturbation amplitude. Similarly, the unstable manifold initial conditions are given by

$$\mathbf{x}_{0,i}^{u,\pm} = \mathbf{x}^*(t_i) \pm \varepsilon \frac{\mathbf{v}_u(t_i)}{\|\mathbf{v}_u(t_i)\|}. \tag{3.24}$$

The signs $\pm$ account for the two branches associated with each eigendirection. Therefore, for each sampled point of the periodic orbit, two stable branches and two unstable branches can be generated.

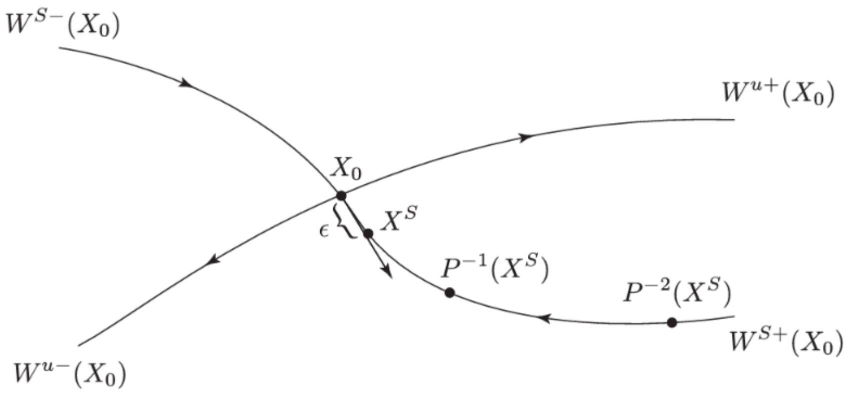

Figure 3.4: Local generation of stable and unstable manifold branches by perturbing a point on the periodic orbit along the corresponding eigendirections. [Reproduced from Koon et al. [2]]

After the perturbation is applied, the resulting states are propagated using the full nonlinear CR3BP dynamics. Unstable manifolds are generated by propagating the perturbed states forward in time, since they naturally depart from the periodic orbit. Stable manifolds, instead, are commonly generated by propagating the perturbed states backward in time. When these backward-propagated arcs are reversed, they represent trajectories that approach the periodic orbit in forward time.

Repeating this procedure for several sampled points along the periodic orbit produces a discrete approximation of the invariant manifold structure. The quality of this approximation depends on the number of sampled points, the perturbation amplitude ε, the propagation time, and the numerical accuracy of the integration scheme.

3.9 Generation of the Stable Manifold Library

For the control strategy developed in the following chapters, the stable invariant manifolds are generated offline and stored in a dedicated library. This avoids computing manifold trajectories online during the NMPC simulation and provides a set of dynamically consistent candidate references.

The starting point is a periodic orbit around the Earth–Moon L_2 point, sampled at N_p points:

$$\mathcal{O} = \{\mathbf{x}^*(t_i)\}_{i=1}^{N_p} . \tag{3.25}$$

For each sampled point, the local stable eigendirection is computed and used to initialize the two stable manifold branches according to Eq. (3.23). Each perturbed initial condition is then propagated backward in time using the nonlinear CR3BP dynamics. The propagated arcs are reversed before storage, so that each stored trajectory evolves toward the periodic orbit in forward time.

Each stable manifold trajectory is stored as a discrete sequence of states:

$$\mathcal{M}_i^{s,\pm} = \left\{\mathbf{x}_{i,0}^{s,\pm}, \mathbf{x}_{i,1}^{s,\pm}, \ldots, \mathbf{x}_{i,N_m}^{s,\pm}\right\}, \tag{3.26}$$

where N_m is the number of discrete points along the manifold trajectory. The complete stable manifold library is then defined as

$$\mathcal{L}^s = \left\{ \mathcal{M}_i^{s,+}, \mathcal{M}_i^{s,-} \right\}_{i=1}^{N_p}. \tag{3.27}$$

In the numerical implementation, each trajectory is stored together with its branch identifier and the index of the corresponding point on the periodic orbit. This structure allows the controller to access, compare, and select candidate manifold trajectories efficiently during the closed-loop simulation.

A key aspect of the library generation is numerical consistency. The same fourth-order Runge–Kutta discretization used in the NMPC prediction model is also adopted to propagate the manifold trajectories. In this way, the reference trajectories stored in the library are dynamically compatible with the model used by the optimizer, reducing artificial discrepancies between the guidance references and the predicted spacecraft motion.

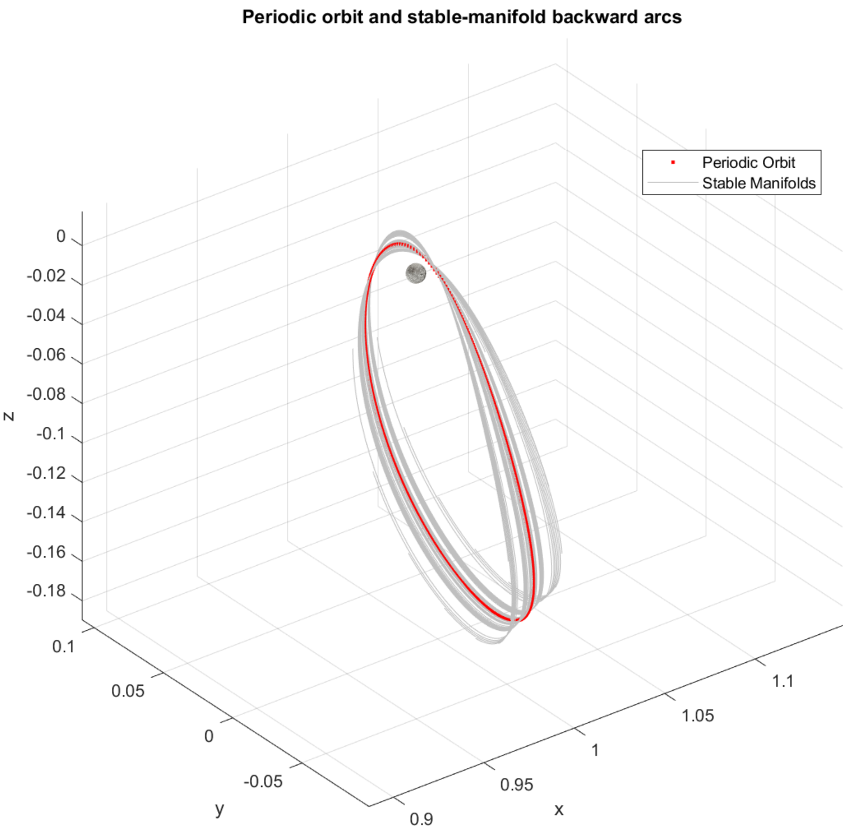

Figure 3.5: Stable manifold trajectories generated from sampled points of a periodic orbit around the Earth–Moon L_2 point. The branchs are obtained by perturbing the orbit along the local stable eigendirections and propagating the resulting states backward in time.

The resulting library contains multiple stable manifold trajectories associated with different points of the base periodic orbit. In the proposed NMPC architecture, these trajectories are used as natural reference candidates. The controller can therefore select, advance, and switch among different manifold branches according to the current spacecraft state, while the artificial reference formulation provides the flexibility required to connect the controlled trajectory to the selected natural path.

3.10 Concluding Remarks

This chapter introduced the invariant manifolds associated with periodic orbits in the CR3BP. Starting from the local linearization of the equations of motion, the State Transition Matrix and the monodromy matrix were used to characterize the stability properties of periodic orbits.

The eigenstructure of the monodromy matrix provides the stable and unstable directions required to initialize manifold trajectories. By perturbing points along the periodic orbit and propagating the resulting states with the full nonlinear dynamics, stable and unstable manifold branches can be numerically constructed.

These structures play a central role in the proposed guidance and control strategy. In the following chapters, stable manifold trajectories are used as dynamically consistent references inside an NMPC framework, allowing the controller to exploit the natural geometry of the CR3BP rather than tracking an arbitrary path.

Chapter 4

Nonlinear Model Predictive Control

4.1 Overview of Model Predictive Control

Model Predictive Control (MPC) is an optimization-based control strategy in which the control action is computed by solving a finite-horizon optimal control problem. At each sampling instant, the future evolution of the system is predicted over a finite time interval, and an optimal sequence of control inputs is obtained by minimizing a prescribed cost function subject to the system dynamics and constraints.

Only the first input of the optimal sequence is applied to the system. At the next sampling instant, the state is measured or estimated again, the prediction horizon is shifted forward, and a new optimization problem is solved. This repeated procedure is known as the receding-horizon strategy and provides feedback to the controller.

The general theory, computational aspects, and design principles of MPC are extensively discussed in [7]. In this thesis, the nonlinear version of MPC is considered, since the spacecraft motion in the EM CR3BP is strongly nonlinear.

4.2 Nonlinear Model Predictive Control Formulation

Nonlinear Model Predictive Control (NMPC) extends the MPC framework to systems described by nonlinear dynamics. Let the continuous-time controlled system be written as

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}), \tag{4.1}$$

where $\mathbf{x} \in \mathbb{R}^n$ is the state vector and $\mathbf{u} \in \mathbb{R}^m$ is the control input.

After discretization with sampling time T_s, the prediction model can be written as

$$\mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k), \qquad k = 0, \dots, N-1, \tag{4.2}$$

where N is the prediction horizon and $\mathbf{F}$ is the discrete-time flow map associated with the continuous dynamics.

At a generic sampling instant t_j, the predicted state and control sequences are

$$\mathbf{X} = \{\mathbf{x}_0, \mathbf{x}_1, \ldots, \mathbf{x}_N\}, \tag{4.3}$$

and

$$\mathbf{U} = \{\mathbf{u}_0, \mathbf{u}_1, \ldots, \mathbf{u}_{N-1}\}. \tag{4.4}$$

The finite-horizon NMPC problem can then be written as

$$\begin{aligned} \min_{\mathbf{x}_0,\ldots,\mathbf{x}_N,\, \mathbf{u}_0,\ldots,\mathbf{u}_{N-1}} \quad & J \\ \text{subject to} \quad & \mathbf{x}_0 = \mathbf{x}(t_j), \\ & \mathbf{x}_{k+1} = \mathbf{F}(\mathbf{x}_k, \mathbf{u}_k), \qquad k = 0, \ldots, N-1, \\ & \mathbf{u}_{\min} \leq \mathbf{u}_k \leq \mathbf{u}_{\max}, \qquad k = 0, \ldots, N-1. \end{aligned} \tag{4.5}$$

The first constraint imposes consistency between the predicted trajectory and the current state of the system. The second constraint enforces the nonlinear discrete-time dynamics along the prediction horizon. The last constraint represents the admissible bounds on the control input.

After discretization, the finite-horizon optimal control problem is transcribed into a Nonlinear Programming Problem (NLP), in which the predicted states and controls are treated as optimization variables.

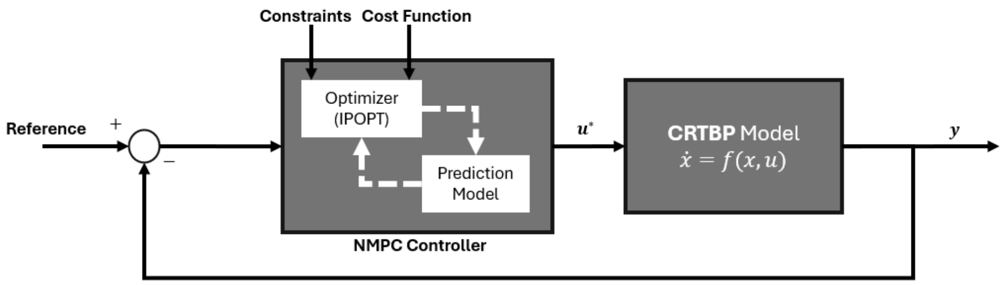

Figure 4.1: Closed-loop NMPC architecture. At each sampling instant, the current state is used to solve a finite-horizon optimal control problem. The first optimal input u^* is applied to the controlled dynamics, and the resulting state is fed back at the next sampling instant.

4.2.1 Terminal Ingredients and Stability

In finite-horizon MPC, the prediction is performed over a limited number of future steps. As a consequence, the behavior of the system beyond the prediction horizon is not optimized directly. For this reason, terminal ingredients are often introduced to influence the final predicted state and to improve the closed-loop behavior of the controller [7,8].

A common possibility is to include a terminal cost, which penalizes the deviation of the final predicted state from a desired terminal condition:

$$\|\mathbf{x}_N - \mathbf{x}_{ref,N}\|_P^2, \tag{4.6}$$

where P is a positive semidefinite terminal weighting matrix. This is a soft formulation, since the terminal state is not required to satisfy a prescribed condition exactly. In standard MPC formulations, P can be chosen as the solution of a Discrete Algebraic Riccati Equation (DARE) associated with a linearized discrete-time model around the reference trajectory.

Alternatively, a terminal constraint can be imposed. In this case, the final predicted state is required to belong to a prescribed terminal set,

$$\mathbf{x}_N \in \mathcal{X}_f, \tag{4.7}$$

or, in a stricter formulation, to satisfy a terminal equality condition,

$$\mathbf{x}_N = \mathbf{x}_{ref,N}. \tag{4.8}$$

This is a hard formulation, since the terminal condition must be satisfied by the optimizer.

Terminal costs and terminal constraints are commonly used to improve the stability properties of finite-horizon MPC schemes. A terminal cost provides a soft approximation of the future cost beyond the prediction horizon, while a terminal constraint imposes an explicit requirement on the final predicted state. However, terminal constraints also make the optimization problem more restrictive and may affect feasibility if the prediction horizon is too short, the input bounds are too tight, or the terminal condition is too demanding.

In the formulation adopted in this thesis, terminal behavior is handled through a terminal constraint. This choice is motivated by the need to enforce a prescribed connection at the end of the prediction horizon. The specific form of this constraint is introduced in the following chapters, where the artificial reference formulation is developed.

4.3 Cost Function and Constraints

The cost function defines the desired behavior of the closed-loop system. In a standard trajectory-tracking formulation, a reference trajectory

$$\mathbf{x}_{ref,k}, \qquad k = 0, \dots, N, \tag{4.9}$$

is prescribed over the prediction horizon. The cost function commonly penalizes the tracking error, the control effort, and, if required, variations of the control input.

A standard quadratic tracking cost is

$$J = \sum_{k=0}^{N-1} \left[\|\mathbf{x}_k - \mathbf{x}_{ref,k}\|_Q^2 + \|\mathbf{u}_k\|_R^2 \right] + \|\mathbf{x}_N - \mathbf{x}_{ref,N}\|_P^2 , \tag{4.10}$$

where

$$\|\mathbf{a}\|_Q^2 = \mathbf{a}^T Q \mathbf{a} \tag{4.11}$$

denotes a weighted quadratic norm.

The matrices Q and P weight the tracking error along the horizon and at the terminal point, respectively, while the matrix R penalizes the magnitude of the control input.

The input constraints

$$\mathbf{u}_{\min} \leq \mathbf{u}_k \leq \mathbf{u}_{\max} \tag{4.12}$$

are particularly important in spacecraft applications, where the available thrust or acceleration is limited. One of the main advantages of NMPC is that such constraints can be included directly in the optimization problem.

4.4 Receding-Horizon Implementation

Once the optimal control sequence has been computed,

$$\mathbf{U}^* = \left\{\mathbf{u}_0^*, \mathbf{u}_1^*, \ldots, \mathbf{u}_{N-1}^*\right\}, \tag{4.13}$$

only the first input $\mathbf{u}_0^*$ is applied to the system. The state is then propagated over one sampling interval, and the optimization problem is solved again from the new state.

This receding-horizon strategy is the mechanism that turns the finite-horizon optimization problem into a feedback controller. The controller continuously updates its prediction according to the current state, making it possible to react to disturbances, numerical deviations, and model mismatches.

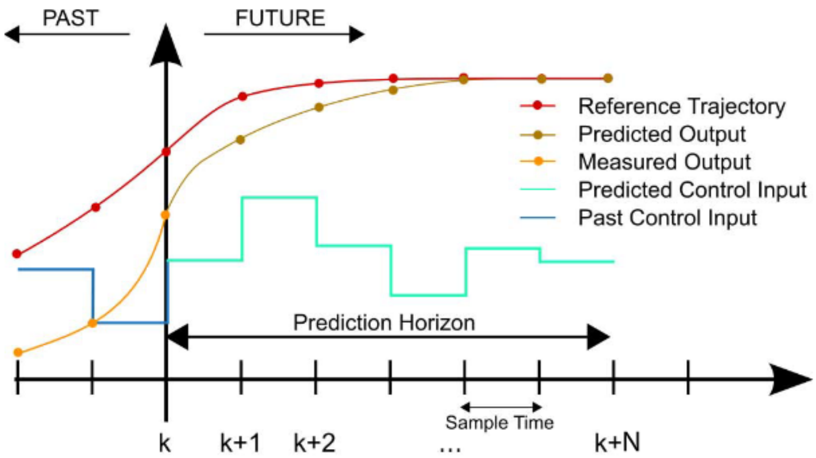

Figure 4.2: Receding-horizon concept of Model Predictive Control. [Reproduced from Martin Behrendt, Wikimedia Commons]

4.5 Application to the CR3BP Guidance Problem

In the present thesis, the NMPC framework is applied to the spacecraft guidance problem in the EM CR3BP. The state vector is

$$\mathbf{x} = \begin{bmatrix} x & y & z & \dot{x} & \dot{y} & \dot{z} \end{bmatrix}^T \in \mathbb{R}^6, \tag{4.14}$$

and the control input is the acceleration vector

$$\mathbf{u} = \begin{bmatrix} u_x & u_y & u_z \end{bmatrix}^T \in \mathbb{R}^3. \tag{4.15}$$

The continuous-time prediction model is given by the controlled CR3BP equations of motion,

$$\dot{\mathbf{x}} = \mathbf{f}_{CR3BP}(\mathbf{x}, \mathbf{u}), \tag{4.16}$$

where the control input acts directly on the acceleration components.

For numerical implementation, the continuous-time dynamics are discretized using the fourth-order Runge–Kutta scheme introduced in Section 2.9. The resulting prediction model is

$$\mathbf{x}_{k+1} = \mathbf{F}_d(\mathbf{x}_k, \mathbf{u}_k). \tag{4.17}$$

NMPC is suitable for this problem for three main reasons. First, it allows the nonlinear CR3BP dynamics to be included directly in the prediction model. Second, it handles bounds on the control input explicitly. Third, the finite prediction horizon allows the controller to account for the future evolution of the spacecraft, which is essential near unstable periodic orbits.

The use of NMPC in lunar and cislunar guidance problems has also been investigated in recent works. For example, Pagone et al. [9] applied an NMPC strategy combined with Pontryagin's principle to autonomous lunar rendezvous trajectory planning and control in non-Keplerian lunar orbits. This provides an additional motivation for adopting predictive nonlinear control techniques in the same multi-body dynamical environment considered in this thesis.

However, a standard tracking formulation is not sufficient by itself. If the reference trajectory is fixed and time-parametrized, the controller may try to synchronize the spacecraft with a specific point of an unstable trajectory. This can lead to persistent control activity and may prevent the exploitation of nearby dynamically consistent paths, such as stable manifold trajectories.

For this reason, the NMPC formulation introduced in this chapter will be extended in the following chapters through the use of an artificial reference and a manifold-based reference selection strategy.

4.6 Numerical Implementation with CasADi

The NMPC problem leads to a nonlinear constrained optimization problem. In this thesis, the resulting NLP is implemented using CasADi, a software framework for nonlinear optimization and optimal control [10].

CasADi provides symbolic modeling capabilities and automatic differentiation, which are useful for constructing the nonlinear dynamics, the RK4 prediction model, the cost function, and the constraint equations. The symbolic expressions are then assembled into an NLP and passed to a numerical solver.

In the implementation, the predicted states and control inputs over the horizon are introduced as optimization variables. The initial condition and the RK4 dynamics are imposed through equality constraints, while the control bounds are imposed as inequality or bound constraints. At each sampling instant, the NLP is solved, the first optimal control input is applied to the spacecraft, and the horizon is shifted forward.

4.7 Concluding Remarks

This chapter introduced the NMPC framework used as the basis for the proposed guidance strategy. The main elements of NMPC were presented: finite-horizon prediction, nonlinear dynamics, cost function design, input constraints, and receding-horizon implementation.

The formulation was then connected to the CR3BP guidance problem, where the nonlinear dynamics and the instability of periodic orbits make predictive constrained control particularly suitable. Finally, the numerical implementation using CasADi was introduced.

The next chapter extends this standard NMPC formulation by introducing an artificial reference, which provides additional flexibility in connecting the controlled spacecraft trajectory to dynamically consistent manifold references.

Chapter 5

Artificial Reference NMPC Formulation

5.1 Introduction

The previous chapter introduced the standard NMPC framework and discussed its application to the CR3BP guidance problem. Although NMPC provides a suitable structure for handling nonlinear dynamics and input constraints, a direct tracking formulation may still be too restrictive when the reference trajectory is unstable, time-dependent, or not well aligned with the current spacecraft state.

In standard tracking NMPC, the predicted spacecraft trajectory is directly penalized with respect to a prescribed reference. This means that the optimizer is required to synchronize the predicted state with a fixed time-parametrized trajectory over the prediction horizon. In the CR3BP, this requirement can be demanding, especially when the reference is associated with unstable periodic or manifold trajectories.

To introduce additional flexibility, this chapter presents an artificial reference formulation. The artificial reference is an auxiliary trajectory introduced inside the NMPC problem. It provides a flexible connection between the real spacecraft motion and the selected natural reference. Instead of forcing the spacecraft to track the natural reference directly at every prediction step, the controller drives the real trajectory toward the artificial trajectory, while the artificial trajectory is itself attracted toward the natural reference.

This two-level structure can be interpreted as a staged connection:

$$\mathbf{X} \;\longrightarrow\; \mathbf{X}_s \;\longrightarrow\; \hat{\mathbf{X}}_s, \tag{5.1}$$

where $\mathbf{X}$ denotes the predicted controlled spacecraft trajectory, $\mathbf{X}_s$ is the optimized artificial reference trajectory, and $\hat{\mathbf{X}}_s$ represents the desired natural reference trajectory. In this way, the artificial trajectory acts as an intermediate bridge between the real spacecraft motion and the selected natural reference.

This formulation provides a more flexible structure for connecting the controlled spacecraft motion to dynamically meaningful reference trajectories. The complete integration with the manifold library and the reference selection logic is developed in the following chapter.

5.2 Motivation for Artificial References

In constrained predictive control problems, the desired reference may not always be directly reachable from the current state while satisfying the system dynamics and the imposed constraints. In such cases, forcing the controller to track the prescribed reference exactly can lead to excessive control effort, poor transient behavior, or even loss of feasibility.

Artificial-reference formulations address this issue by introducing an auxiliary reference inside the optimization problem. Instead of treating the desired reference as the only admissible target, the controller optimizes an intermediate reference that is compatible with the current state, the dynamics, and the constraints. This artificial reference is then driven toward the desired one through suitable cost terms.

This idea has been developed in the context of MPC for tracking constrained systems. For constrained linear systems, Limón et al. [3] proposed an MPC formulation for tracking piecewise constant references in which an artificial steady state is optimized together with the input sequence. The artificial steady state is chosen to be admissible, while an offset cost penalizes its deviation from the desired reference.

The basic mechanism is illustrated in Figure 5.1, which was generated following the artificial-reference tracking framework proposed in [3]. The desired reference is not imposed as a rigid target to be followed directly. Instead, an artificial reference is optimized inside the control problem and evolves toward the desired target while remaining compatible with the imposed constraints. The controlled trajectory is then guided toward this artificial reference.

The initial transient illustrates why an artificial reference is useful. Since the real system is constrained by its dynamics and by bounded inputs, the output cannot immediately follow a sudden change in the desired reference. The artificial reference evolves toward the target more gradually and provides an intermediate trajectory that the real system can track more easily.

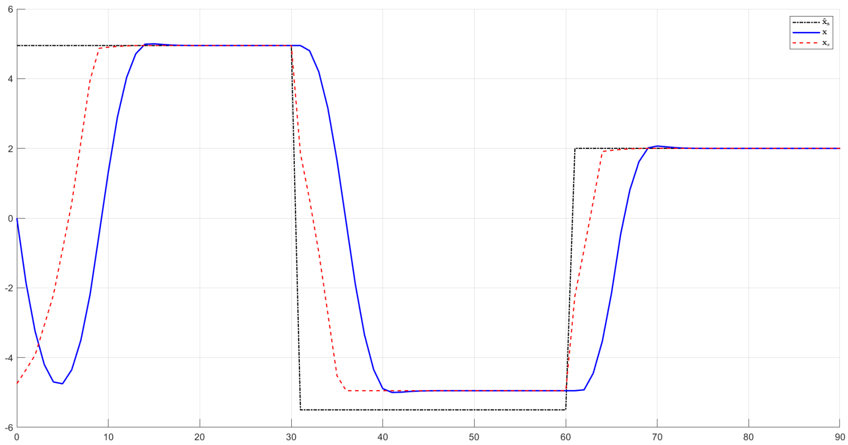

Figure 5.1: Conceptual illustration of the artificial-reference tracking structure. Following the framework of Limón et al. [3], the real trajectory X (solid blue line) is driven toward the artificial trajectory X_s (red dashed line), while the artificial trajectory is attracted toward the desired natural reference $\hat{X}_s$ (dash-dotted black line).

The same concept was later extended to constrained nonlinear systems in [11]. A recent overview of MPC for tracking using artificial references, including theoretical properties, extensions, and implementation aspects, is provided in [12].

The formulation adopted in this thesis is inspired by this general philosophy, but it is adapted to the trajectory-guidance problem in the CR3BP. In the classical tracking formulations mentioned above, the artificial reference is usually a steady state or a piecewise constant reference. In the present case, the desired reference is instead a time-varying natural trajectory, associated with a stable manifold of a periodic orbit. Therefore, the auxiliary reference is introduced as an artificial trajectory rather than as a constant set-point.

This distinction is important because the difficulty is not limited to reaching an admissible steady state. The controller must also connect the current spacecraft motion to a selected natural trajectory whose phase, index, or branch may not be immediately compatible with the current state. The artificial trajectory provides the additional degrees of freedom needed to construct such a connection while preserving the constrained predictive-control structure.

Thus, the artificial reference is used here as a flexible intermediate trajectory that reduces the rigidity of direct tracking and prepares the integration with the manifold-based guidance strategy developed in the following chapter.

5.3 Artificial Trajectory Definition

The artificial reference is represented by an auxiliary state sequence

$$\mathbf{X}_s = \{\mathbf{x}_{s,0}, \mathbf{x}_{s,1}, \dots, \mathbf{x}_{s,N}\}, \tag{5.2}$$

where $\mathbf{x}_{s,k} \in \mathbb{R}^6$ denotes the artificial state at prediction step k.

The artificial trajectory is propagated through its own discrete-time dynamics:

$$\mathbf{x}_{s,k+1} = \mathbf{F}_d\left(\mathbf{x}_{s,k}, \mathbf{u}_{s,k}\right), \qquad k = 0, \ldots, N-1, \tag{5.3}$$

where $\mathbf{u}_{s,k}$ is the artificial control input. The same RK4 discretization used for the real spacecraft prediction is adopted for the artificial trajectory. This ensures that both trajectories are propagated consistently with the same discrete-time CR3BP model.

The artificial input $\mathbf{u}_{s,k}$ does not represent a physical command applied to the spacecraft. Rather, it is an optimization variable that gives the artificial trajectory additional flexibility. By penalizing $\mathbf{u}_{s,k}$ in the cost function and by imposing suitable bounds, the artificial trajectory is encouraged to remain dynamically consistent with the natural CR3BP flow while still being able to connect the real trajectory to the selected reference.

5.4 Two-Level Tracking Structure

The artificial reference formulation introduces a two-level tracking structure. The first level connects the real controlled trajectory to the artificial trajectory. The second level connects the artificial trajectory to the natural reference.

Let $\mathbf{x}_k$ denote the predicted real spacecraft state and let $\mathbf{x}_{s,k}$ denote the artificial state. The real-to-artificial tracking error is

$$\mathbf{e}_{x,k} = \mathbf{x}_k - \mathbf{x}_{s,k}. \tag{5.4}$$

Let $\mathbf{x}_{ref,k}$ denote the selected natural reference trajectory. The artificial-to-reference error is

$$\mathbf{e}_{s,k} = \mathbf{x}_{s,k} - \mathbf{x}_{ref,k}. \tag{5.5}$$

The first error term drives the controlled spacecraft toward the artificial trajectory. The second error term attracts the artificial trajectory toward the natural reference. Therefore, the optimizer is not forced to align the real trajectory directly with the natural reference at every step. Instead, it can construct an intermediate dynamically consistent bridge.

This structure is especially useful when the natural reference is not immediately reachable from the current spacecraft state or when direct synchronization would require excessive control effort.

5.5 Artificial Reference NMPC Problem

The artificial-reference NMPC problem includes both the real and artificial trajectories as optimization variables. The real trajectory satisfies

$$\mathbf{x}_{k+1} = \mathbf{F}_d\left(\mathbf{x}_k, \mathbf{u}_k\right), \qquad k = 0, \ldots, N-1, \tag{5.6}$$

while the artificial trajectory satisfies Eq. (5.3).

The optimization variables are therefore

$$\mathbf{X}, \qquad \mathbf{U}, \qquad \mathbf{X}_s, \qquad \mathbf{U}_s, \tag{5.7}$$

where

$$\mathbf{U}_s = \{\mathbf{u}_{s,0}, \mathbf{u}_{s,1}, \dots, \mathbf{u}_{s,N-1}\}\,. \tag{5.8}$$

A compact formulation is

$$\begin{aligned}
\min_{\mathbf{X},\mathbf{U},\mathbf{X}_s,\mathbf{U}_s} \quad & J \\
\text{subject to} \quad & \mathbf{x}_0 = \mathbf{x}(t_j), \\
& \mathbf{x}_{k+1} = \mathbf{F}_d(\mathbf{x}_k, \mathbf{u}_k), \qquad k = 0, \dots, N-1, \\
& \mathbf{x}_{s,k+1} = \mathbf{F}_d(\mathbf{x}_{s,k}, \mathbf{u}_{s,k}), \qquad k = 0, \dots, N-1, \\
& \mathbf{u}_{\min} \le \mathbf{u}_k \le \mathbf{u}_{\max}, \qquad k = 0, \dots, N-1, \\
& \mathbf{u}_{s,\min} \le \mathbf{u}_{s,k} \le \mathbf{u}_{s,\max}, \qquad k = 0, \dots, N-1, \\
& \mathbf{x}_N = \mathbf{x}_{s,N}.
\end{aligned} \tag{5.9}$$

The last condition is the terminal constraint. It enforces the connection between the real and artificial trajectories at the end of the prediction horizon. This constraint prevents the optimizer from keeping the real and artificial trajectories separated indefinitely and gives the bridge structure a precise endpoint.

5.6 Cost Function

The cost function is designed to couple the real controlled trajectory, the artificial trajectory, and the selected natural reference. It contains three main tracking contributions: the distance between the real and artificial trajectories, the distance between the artificial trajectory and the natural reference, and the mismatch between the real and artificial control inputs.

A representative form of the stage cost is

$$\begin{aligned}
\ell_k = {} & \|\mathbf{x}_k - \mathbf{x}_{s,k}\|_Q^2 + \|\mathbf{x}_{s,k} - \hat{\mathbf{x}}_k\|_{Q_s}^2 + \|\mathbf{u}_k - \mathbf{u}_{s,k}\|_R^2 \\
& + \|\Delta\mathbf{u}_k\|_{R_{\Delta u}}^2 + \|\Delta\mathbf{u}_{s,k}\|_{R_{\Delta u_s}}^2 .
\end{aligned} \tag{5.10}$$

The first term drives the real trajectory toward the artificial trajectory. The second term attracts the artificial trajectory toward the selected natural reference. The third term penalizes the mismatch between the real and artificial control inputs, encouraging the two trajectories to remain dynamically compatible.

The terms involving $\Delta\mathbf{u}_k$ and $\Delta\mathbf{u}_{s,k}$ regularize the real and artificial control profiles, reducing abrupt variations between consecutive prediction steps.

The terminal contribution is defined as

$$\ell_N = \|\mathbf{x}_{s,N} - \hat{\mathbf{x}}_N\|_{P_s}^2 . \tag{5.11}$$

The complete cost function is then

$$J = \sum_{k=0}^{N-1} \ell_k + \ell_N. \tag{5.12}$$

5.7 Interpretation of the Terminal Constraint

The terminal constraint

$$\mathbf{x}_N = \mathbf{x}_{s,N} \tag{5.13}$$

is central to the artificial-reference formulation. As discussed in Section 4.2.1, terminal constraints are commonly used in finite-horizon MPC to impose a prescribed behavior at the end of the prediction horizon and to influence the closed-loop behavior of the controller.

In the present formulation, the terminal constraint does not force the predicted spacecraft trajectory to reach the natural reference directly. Instead, it enforces that, at the end of the horizon, the controlled trajectory connects to the artificial trajectory. Therefore, the artificial trajectory acts as a bridge that the real spacecraft must reach within the prediction horizon.

This idea is consistent with artificial-reference MPC formulations, where the terminal condition is imposed with respect to an optimized auxiliary reference rather than directly on the desired target [3, 11, 12]. The purpose is to preserve the role of the terminal constraint while avoiding the excessive rigidity of a direct terminal condition on the desired reference.

Indeed, imposing the terminal constraint directly on the natural reference,

$$\mathbf{x}_N = \hat{\mathbf{x}}_N, \tag{5.14}$$

would be significantly more restrictive. Such a condition would require the real trajectory to match a specific point of the natural reference at the end of the horizon. In the CR3BP, this may be difficult because the selected manifold trajectory may not be immediately synchronized with the current spacecraft state, and the available control input is bounded.

By contrast, the artificial trajectory is optimized together with the real control sequence and the artificial input sequence. This provides additional flexibility: the real trajectory is forced to connect to the artificial one, while the artificial trajectory is simultaneously attracted toward the natural reference through the cost term

$$\|\mathbf{x}_{s,k} - \hat{\mathbf{x}}_k\|^2_{Q_s}. \tag{5.15}$$

As a result, the terminal constraint gives the prediction a precise endpoint, while the artificial reference allows the optimizer to shape the bridge in a way that remains compatible with the nonlinear dynamics and the input constraints. This is particularly useful for manifold-based guidance, where the desired reference is dynamically meaningful but may not be directly reachable as a rigid terminal target.

5.8 Role in the Manifold-Based Guidance Strategy

The formulation presented in this chapter is independent of the specific source of the natural reference. In the following chapter, the natural reference $\mathbf{x}_{ref,k}$ is selected from the stable manifold library. In that context, the artificial reference provides the mechanism that connects the real controlled trajectory to a selected manifold branch.

This is important because stable manifold trajectories are dynamically meaningful but not necessarily synchronized with the current spacecraft state. A direct tracking formulation would require the spacecraft to follow a prescribed manifold trajectory with a fixed phase and index. The artificial reference formulation relaxes this requirement by allowing an optimized intermediate trajectory.

Therefore, the artificial reference is the element that links the predictive control problem to the manifold-based guidance strategy. The manifold library provides natural candidate references, while the artificial trajectory provides the flexibility required to connect to them.

5.9 Concluding Remarks

This chapter introduced the artificial reference NMPC formulation. Starting from the limitations of direct tracking in the CR3BP, an auxiliary trajectory was introduced as an intermediate bridge between the real controlled spacecraft trajectory and the selected natural reference.

The resulting formulation includes both real and artificial dynamics, a two-level tracking cost, input bounds for both real and artificial controls, and a terminal constraint connecting the real and artificial trajectories at the end of the horizon. This structure preserves the predictive and constrained nature of NMPC while reducing the rigidity of direct reference tracking.

The next chapter integrates this artificial-reference formulation with the stable manifold library, leading to the complete manifold-based NMPC strategy.

Chapter 6

Manifold-Based NMPC Guidance Strategy

6.1 Introduction

The previous chapters introduced the two main components of the proposed guidance strategy: the stable manifold library and the artificial-reference NMPC formulation. The purpose of this chapter is to describe how these elements are combined within the complete closed-loop architecture.

At each sampling instant, the current spacecraft state is compared with the trajectories stored in the manifold library. A suitable manifold state is selected as a reference seed and propagated forward under the uncontrolled CR3BP dynamics to generate the natural reference over the NMPC prediction horizon. The resulting reference is then supplied to the artificial-reference NMPC problem.

The architecture also includes mechanisms for advancing the active seed along the selected trajectory, switching among different manifold candidates, and activating coasting phases when the uncontrolled motion remains sufficiently close to the active natural path.

For clarity, ℓ denotes the closed-loop sampling instant, k denotes the prediction step within the NMPC horizon, j identifies a stored manifold trajectory, and q identifies a sample along that trajectory.

6.2 Stable Manifold Candidates

The stable manifold library generated in Chapter 3 contains a collection of trajectories associated with different points of the selected halo orbit. In the present implementation, three candidate trajectories are extracted from the library and denoted by

$$\mathcal{M}_A^s, \qquad \mathcal{M}_B^s, \qquad \mathcal{M}_C^s. \tag{6.1}$$

More generally, the set of candidate trajectories is written as

$$\mathcal{L}_c^s = \left\{ \mathcal{M}_j^s \right\}_{j=1}^{N_{\text{traj}}}, \tag{6.2}$$

where N_{traj} is the number of manifold trajectories considered by the online guidance algorithm.

Each candidate trajectory is stored as an ordered sequence of states,

$$\mathcal{M}_j^s = \left\{ \hat{\mathbf{x}}_{j,1}, \hat{\mathbf{x}}_{j,2}, \dots, \hat{\mathbf{x}}_{j,N_j} \right\}, \tag{6.3}$$

where

$$\hat{\mathbf{x}}_{j,q} = \begin{bmatrix} \hat{\mathbf{r}}_{j,q} \\ \hat{\mathbf{v}}_{j,q} \end{bmatrix} \in \mathbb{R}^6. \tag{6.4}$$

The original stable arcs are generated by backward propagation from the halo orbit. Before being used by the controller, their ordering is reversed so that increasing q corresponds to forward evolution from the outer part of the manifold toward the periodic orbit.

Therefore, a pair (j, q) identifies both the active manifold trajectory and the current seed point along it.

6.3 Initial Reference Selection

The simulation is initialized from one of the candidate trajectories. Let the initial spacecraft state be denoted by

$$\mathbf{x}_0 = \begin{bmatrix} \mathbf{r}_0 \\ \mathbf{v}_0 \end{bmatrix}. \tag{6.5}$$

The initial active trajectory is selected in advance, while the initial seed index is determined by comparing $\mathbf{x}_0$ with all points of that trajectory.

A weighted state-space distance is used:

$$d_{j,q} = \left\| W_{\text{search}} \odot (\hat{\mathbf{x}}_{j,q} - \mathbf{x}_0) \right\|_2^2, \tag{6.6}$$

where $\odot$ denotes component-wise multiplication and $W_{\text{search}} \in \mathbb{R}^6$ contains the position and velocity weights.

The initial seed index is then obtained as

$$q_0^\star = \arg\min_q d_{j_0,q}, \tag{6.7}$$

where j_0 is the initially active trajectory.

The corresponding stored state,

$$\hat{\mathbf{x}}_{\text{seed},0} = \hat{\mathbf{x}}_{j_0,q_0^\star}, \tag{6.8}$$

is used to construct the first natural reference horizon.

6.4 Natural Reference Generation

The manifold library is used to select the initial state of the natural reference, rather than to provide the complete prediction segment directly.

At the ℓ-th closed-loop step, let

$$\hat{\mathbf{x}}_{\text{seed},\ell} = \hat{\mathbf{x}}_{j_\ell,q_\ell} \tag{6.9}$$

be the active seed state selected from the current manifold trajectory.

The natural reference is initialized as

$$\hat{\mathbf{x}}_{0|\ell} = \hat{\mathbf{x}}_{\text{seed},\ell}, \tag{6.10}$$

and propagated over the prediction horizon using the uncontrolled discrete-time CR3BP dynamics:

$$\hat{\mathbf{x}}_{k+1|\ell} = F_d\left(\hat{\mathbf{x}}_{k|\ell}, \mathbf{0}\right), \qquad k = 0, \ldots, N-1. \tag{6.11}$$

The complete natural reference supplied to the optimizer is therefore

$$\hat{X}_\ell = \left\{\hat{\mathbf{x}}_{0|\ell}, \hat{\mathbf{x}}_{1|\ell}, \ldots, \hat{\mathbf{x}}_{N|\ell}\right\}. \tag{6.12}$$

This construction ensures that the reference satisfies the same RK4-discretized CR3BP model used inside the NMPC problem. It also avoids possible inconsistencies caused by directly extracting a sequence of library points whose sampling may not exactly coincide with the prediction discretization.

6.5 Integration with the Artificial-Reference NMPC

The reference horizon $\hat{X}_\ell$ is passed to the artificial-reference NMPC problem introduced in Chapter 5.

At the ℓ-th sampling instant, the real predicted trajectories are

$$X_\ell = \left\{\mathbf{x}_{0|\ell}, \ldots, \mathbf{x}_{N|\ell}\right\}, \qquad U_\ell = \left\{\mathbf{u}_{0|\ell}, \ldots, \mathbf{u}_{N-1|\ell}\right\}, \tag{6.13}$$

while the artificial trajectories are

$$X_{s,\ell} = \left\{\mathbf{x}_{s,0|\ell}, \ldots, \mathbf{x}_{s,N|\ell}\right\}, \qquad U_{s,\ell} = \left\{\mathbf{u}_{s,0|\ell}, \ldots, \mathbf{u}_{s,N-1|\ell}\right\}. \tag{6.14}$$

Both trajectories satisfy the same discrete-time controlled CR3BP model:

$$\mathbf{x}_{k+1|\ell} = F_d\left(\mathbf{x}_{k|\ell}, \mathbf{u}_{k|\ell}\right), \tag{6.15}$$

$$\mathbf{x}_{s,k+1|\ell} = F_d\left(\mathbf{x}_{s,k|\ell}, \mathbf{u}_{s,k|\ell}\right), \tag{6.16}$$

for $k = 0, \ldots, N-1$.

The real prediction is initialized from the current closed-loop state:

$$\mathbf{x}_{0|\ell} = \mathbf{x}_\ell. \tag{6.17}$$

The initial artificial state remains an optimization variable. The artificial trajectory can therefore adjust between the current spacecraft state and the selected natural reference.

The real and artificial predictions are connected by the terminal equality constraint

$$\mathbf{x}_{N|\ell} = \mathbf{x}_{s,N|\ell}. \tag{6.18}$$

After solving the nonlinear programming problem, only the first real control input

$$\mathbf{u}^{\star}_{0|\ell} \tag{6.19}$$

is applied. The spacecraft state is then propagated by one sampling interval:

$$\mathbf{x}_{\ell+1} = F_d\left(\mathbf{x}_\ell, \mathbf{u}^{\star}_{0|\ell}\right). \tag{6.20}$$

6.6 Coasting Pre-Check

Before solving the NMPC problem, the algorithm evaluates whether the spacecraft can safely evolve for one sampling interval without control.

A one-step uncontrolled prediction is computed as

$$\mathbf{x}^{\text{coast}}_{\ell+1} = F_d\left(\mathbf{x}_\ell, \mathbf{0}\right). \tag{6.21}$$

This predicted state is compared with a limited forward window of the active manifold trajectory. If q_ℓ is the current seed index, the admissible indices are

$$\mathcal{I}^{\text{coast}}_\ell = \left\{q_\ell + \Delta q_{\min}, \ldots, q_\ell + \Delta q_{\max}\right\}, \tag{6.22}$$

with the upper bound truncated at the end of the stored trajectory.

The best candidate in this interval is

$$q^{\text{coast}}_\ell = \arg\min_{q \in \mathcal{I}^{\text{coast}}_\ell} \left\| W_{\text{advance}} \odot \left(\hat{\mathbf{x}}_{j_\ell,q} - \mathbf{x}^{\text{coast}}_{\ell+1}\right)\right\|_2^2. \tag{6.23}$$

The corresponding position and velocity errors are

$$e^{\text{coast}}_{r,\ell} = \left\|\hat{\mathbf{r}}_{j_\ell,q^{\text{coast}}_\ell} - \mathbf{r}^{\text{coast}}_{\ell+1}\right\|, \tag{6.24}$$

$$e^{\text{coast}}_{v,\ell} = \left\|\hat{\mathbf{v}}_{j_\ell,q^{\text{coast}}_\ell} - \mathbf{v}^{\text{coast}}_{\ell+1}\right\|. \tag{6.25}$$

Coasting is activated when

$$e^{\text{coast}}_{r,\ell} < \varepsilon_r, \qquad e^{\text{coast}}_{v,\ell} < \varepsilon_v. \tag{6.26}$$

If these conditions are satisfied, the applied real and artificial inputs are set to zero:

$$\mathbf{u}_\ell = \mathbf{0}, \qquad \mathbf{u}_{s,\ell} = \mathbf{0}. \tag{6.27}$$

Otherwise, the artificial-reference NMPC problem is solved normally.

The coasting decision is therefore based on the predicted uncontrolled next state rather than only on the current distance from the manifold. This provides a more meaningful indication of whether the natural dynamics will remain compatible with the selected path during the following interval.

6.7 Reference Update on the Active Trajectory

After the next real state $\mathbf{x}_{\ell+1}$ has been obtained, either through NMPC or coasting, the reference seed is updated.

For the currently active trajectory j_ℓ, only a forward index window is considered:

$$\mathcal{I}_\ell^{\mathrm{curr}} = \{q_\ell + \Delta q_{\min}, \dots, q_\ell + \Delta q_{\max}\} . \tag{6.28}$$

The updated seed index on the current trajectory is

$$q^\star_{\mathrm{curr}} = \arg \min_{q \in \mathcal{I}_\ell^{\mathrm{curr}}} \| W_{\mathrm{advance}} \odot (\hat{\mathbf{x}}_{j_\ell, q} - \mathbf{x}_{\ell+1}) \|_2^2 . \tag{6.29}$$

The corresponding score is denoted by

$$S_{\mathrm{curr}}. \tag{6.30}$$

Restricting the search to a forward window prevents backward motion of the reference and limits discontinuous jumps along the active manifold.

6.8 Search Over Alternative Trajectories

The trajectories not currently active are evaluated as switching candidates. In the present implementation, each alternative trajectory is searched globally.

For a candidate trajectory $j \neq j_\ell$, the best index is

$$q^\star_j = \arg \min_q \| W_{\mathrm{advance}} \odot (\hat{\mathbf{x}}_{j,q} - \mathbf{x}_{\ell+1}) \|_2^2 . \tag{6.31}$$

The best alternative is then

$$(j^\star_{\mathrm{alt}}, q^\star_{\mathrm{alt}}) = \arg \min_{j \neq j_\ell} S_j , \tag{6.32}$$

where S_j is the minimum score associated with trajectory j.

The score of the best alternative is denoted by

$$S_{\mathrm{alt}}. \tag{6.33}$$

Thus, the current trajectory is updated locally, while alternative trajectories are evaluated globally. This distinction preserves continuity on the active branch while allowing the controller to identify a substantially better manifold trajectory elsewhere in the library.

6.9 Switching Logic

A candidate trajectory is considered preferable when

$$S_{\mathrm{alt}} < \rho_{\mathrm{sw}} S_{\mathrm{curr}}, \tag{6.34}$$

where

$$0 < \rho_{\mathrm{sw}} < 1 \tag{6.35}$$

is the required improvement ratio.

The condition must remain satisfied for a prescribed number of consecutive closed-loop steps. Let n_{cand} denote the confirmation counter:

$$n_{\mathrm{cand},\ell+1} = \begin{cases} n_{\mathrm{cand},\ell} + 1, & S_{\mathrm{alt}} < \rho_{\mathrm{sw}} S_{\mathrm{curr}}, \\ 0, & \text{otherwise.} \end{cases} \tag{6.36}$$

A switch is performed when

$$n_{\mathrm{cand},\ell+1} \geq N_{\mathrm{hold}}. \tag{6.37}$$

If the switch is accepted, the active trajectory and seed index are updated as

$$j_{\ell+1} = j^{\star}_{\mathrm{alt}}, \qquad q_{\ell+1} = q^{\star}_{\mathrm{alt}}. \tag{6.38}$$

Otherwise, the current trajectory is retained and only its locally updated seed is adopted:

$$j_{\ell+1} = j_{\ell}, \qquad q_{\ell+1} = q^{\star}_{\mathrm{curr}}. \tag{6.39}$$

The hold condition prevents switching caused by temporary score fluctuations and reduces chattering between similar trajectories.

6.10 Warm Starting and Failure Handling

The nonlinear solver is warm-started by shifting the optimal real and artificial trajectories computed at the previous sampling instant.

For the state predictions, the shifted initial guesses are constructed as

$$X^{\mathrm{guess}}_{\ell+1} = \left[\mathbf{x}^{\star}_{1|\ell}, \ldots, \mathbf{x}^{\star}_{N|\ell}, \mathbf{x}^{\star}_{N|\ell}\right], \tag{6.40}$$

with an analogous construction for $X^{\mathrm{guess}}_{s,\ell+1}$, $U^{\mathrm{guess}}_{\ell+1}$, and $U^{\mathrm{guess}}_{s,\ell+1}$.

The previously applied real and artificial inputs are also passed to the next NMPC problem in order to evaluate the first control-rate penalties.

If the nonlinear solver fails, zero real and artificial inputs are applied for the current interval:

$$\mathbf{u}_{\ell} = \mathbf{0}, \qquad \mathbf{u}_{s,\ell} = \mathbf{0}. \tag{6.41}$$

The spacecraft is then propagated using the uncontrolled CR3BP dynamics. This fallback procedure prevents the closed-loop simulation from terminating because of an isolated optimization failure.

6.11 Closed-Loop Procedure

The complete guidance strategy is summarized by the following sequence:

1. Acquire the current spacecraft state $\mathbf{x}_\ell$.
2. Select the active manifold seed $\hat{\mathbf{x}}_{j_\ell,q_\ell}$.
3. Propagate the seed at zero control to construct the natural reference horizon $\hat{X}_\ell$.
4. Compute the one-step uncontrolled prediction $\mathbf{x}_{\ell+1}^{\text{coast}}$ and evaluate the coasting condition.
5. If coasting is admissible, apply zero control.
6. Otherwise, solve the artificial-reference NMPC problem and apply the first optimal real control input.
7. Propagate the spacecraft to obtain $\mathbf{x}_{\ell+1}$.
8. Update the seed on the current trajectory through a local forward search.
9. Evaluate all alternative trajectories through a global search.
10. Apply the switching improvement and hold conditions.
11. Update the active trajectory, seed index, previous inputs, and optimizer initial guesses.
12. Repeat the procedure at the next sampling instant.

The overall information flow can be summarized as

$$\mathcal{L}_c^s \longrightarrow \hat{\mathbf{x}}_{\text{seed},\ell} \longrightarrow \hat{X}_\ell \longrightarrow X_{s,\ell} \longrightarrow X_\ell. \tag{6.42}$$

The manifold library provides the reference seed, the uncontrolled propagation generates the natural reference horizon, the artificial trajectory provides an optimized bridge, and the real predicted trajectory determines the control action applied to the spacecraft.

6.12 Concluding Remarks

This chapter described the integration of the stable manifold candidates with the artificial-reference NMPC formulation.

The manifold library is used to select a reference seed according to the current spacecraft state. The complete natural reference horizon is then generated by propagating this seed under the uncontrolled CR3BP dynamics using the same discrete-time model adopted by the optimizer.

The active seed is advanced through a local forward search, while alternative manifold trajectories are evaluated globally. A switching condition based on a required score improvement and a confirmation time allows the controller to change trajectory without introducing excessive chattering.

The coasting logic evaluates the one-step uncontrolled spacecraft motion before solving the NMPC problem. When this motion remains sufficiently close to the active manifold, zero control is applied and the natural CR3BP dynamics are exploited directly.

The numerical setup and closed-loop results obtained with this architecture are presented in the following chapter.

Chapter 7

Simulation Results

7.1 Introduction

This chapter presents the numerical simulations used to assess the proposed manifold-based artificial-reference NMPC strategy. The analysis considers the complete closed-loop architecture introduced in Chapter 6, including the online selection and advancement of the natural manifold reference, the artificial-reference optimization, trajectory switching, and coasting logic.

The proposed strategy is evaluated in terms of closed-loop trajectory evolution, position and velocity tracking errors, control activity, accumulated Δv, switching behavior, coasting intervals, Jacobi constant evolution, and relative motion in a reference-centered LVLH frame.

A comparison with a conventional NMPC controller is also considered. The conventional formulation directly tracks the target halo orbit without exploiting either the stable manifold library or the artificial-reference trajectory. The comparison is used to evaluate the contribution of the proposed architecture to the overall closed-loop behavior.

7.2 Simulation Setup

The numerical simulations are performed using the nondimensional Earth–Moon CR3BP model introduced in Chapter 2. The same fourth-order Runge–Kutta discretization adopted during the generation of the stable manifold library is used for both the NMPC prediction model and the closed-loop state propagation. This ensures consistency between the stored natural trajectories and the model employed by the optimizer.

The nonlinear optimization problem is formulated in CasADi and solved using IPOPT. At each sampling instant, the solution obtained at the previous step is shifted forward and used as the initial guess for the subsequent optimization problem.

The main numerical simulation parameters are reported in Table 7.1.

Table 7.1: Numerical simulation setup.

Parameter	Symbol	Value
Earth–Moon mass parameter	μ	0.012150585609624
State dimension	n	6
Input dimension	m	3
Prediction horizon	N	50
Closed-loop simulation steps	N_{sim}	2000
Sampling time	T_s	0.0016
RK4 internal substeps	M	5
Nonlinear optimization solver	–	IPOPT
Symbolic optimization framework	–	CasADi

The dimensional results are recovered using the Earth–Moon distance unit

$$DU = 384000 \times 10^3 \text{ m}, \tag{7.1}$$

and the time unit

$$TU = \frac{27.321661 \cdot 24 \cdot 60 \cdot 60}{2\pi} \text{ s}. \tag{7.2}$$

The corresponding acceleration scale and dimensional sampling interval are

$$a_{\text{unit}} = \frac{DU}{TU^2}, \qquad \Delta t_{\text{dim}} = T_s TU. \tag{7.3}$$

These quantities are used to express the applied control acceleration and the accumulated Δv in physical units.

7.3 Controller and Supervisory Parameters

The artificial-reference NMPC formulation includes penalties on the real-to-artificial state error, the artificial-to-natural-reference error, the mismatch between the real and artificial control inputs, and the variation of both input sequences. A terminal penalty is applied to the artificial-to-natural-reference error, while the connection between the real and artificial trajectories is imposed through the terminal equality constraint introduced in Chapter 5.

The controller also includes a supervisory layer responsible for reference advancement, switching among the candidate manifold trajectories, and activation of the coasting mode. The adopted control and supervisory parameters are summarized in Table 7.2.

Table 7.2: Controller, reference-management, switching, and coasting parameters.

Controller Parameter	Symbol	Value
Real input component bound	$u_{\max}$	2×10^{-3}
Artificial input component bound	$u_{s,\max}$	1.8×10^{-3}
Real-to-artificial state weight	Q	diag($8{\times}10^4 I_3$, $10^4 I_3$)
Artificial-to-natural state weight	Q_s	diag($5{\times}10^4 I_3$, $5{\times}10^3 I_3$)
Real/artificial input mismatch weight	R	$10 I_3$
Real input-rate weight	$R_{\Delta u}$	$10^2 I_3$
Artificial input-rate weight	$R_{\Delta u_s}$	$10^2 I_3$
Terminal artificial-reference weight	P_s	$20 Q_s$
Search weighting vector	W_{search}	$[1,1,1,200,200,200]^T$
Advancement weighting vector	W_{advance}	$[1,1,1,200,200,200]^T$
Supervisory Parameter	**Symbol**	**Value**
Minimum index advancement	$\Delta q_{\min}$	0
Maximum index advancement	$\Delta q_{\max}$	5
Switching improvement ratio	ρ_{sw}	0.7
Switching confirmation duration	N_{hold}	3 steps
Position coasting tolerance	ε_r	5×10^{-5}
Velocity coasting tolerance	ε_v	5×10^{-4}

The artificial input bound is selected as

$$u_{s,\max} = 0.9 u_{\max}, \tag{7.4}$$

so that the artificial trajectory is governed by a slightly smaller control authority than the real spacecraft trajectory.

The larger weights assigned to the velocity components in W_{search} and W_{advance} compensate for their smaller numerical scale and reduce the possibility of selecting a manifold point that is close in position but poorly aligned in velocity.

7.4 Initial Conditions and Manifold Candidates

Three stable manifold trajectories are selected from the precomputed library. They correspond to the entries

$$j_A = 200, \qquad j_B = 250, \qquad j_C = 280, \tag{7.5}$$

and, for all three candidates, the first stored stable branch is used:

$$b_A = b_B = b_C = 1. \tag{7.6}$$

The selected candidates are denoted by trajectories A, B, and C. Each stored backward-time manifold arc is reversed before use so that its state ordering corresponds to forward motion from the outer manifold region toward the halo orbit.

The simulation is initialized using trajectory A. Let

$$\hat{\mathbf{x}}_{A,1} \tag{7.7}$$

be its first stored state after reversal. The initial spacecraft state is obtained by adding the perturbation

$$\delta\mathbf{x}_0 = \begin{bmatrix} 10^{-5} \\ -5\times 10^{-5} \\ 5\times 10^{-5} \\ 5\times 10^{-7} \\ -5\times 10^{-7} \\ 5\times 10^{-7} \end{bmatrix} \tag{7.8}$$

to this state:

$$\mathbf{x}_0 = \hat{\mathbf{x}}_{A,1} + \delta\mathbf{x}_0. \tag{7.9}$$

The initial seed index is then obtained by searching for the point of trajectory A that minimizes the weighted state-space distance from $\mathbf{x}_0$.

Table 7.3: Selected stable manifold trajectories and initial condition.

Candidate	Library index	Branch	Role
A	200	1	Initial active trajectory
B	250	1	Alternative candidate
C	280	1	Alternative candidate

7.5 Closed-Loop Trajectory and Reference Evolution

This section analyzes the closed-loop evolution obtained with the proposed manifold-based artificial-reference NMPC strategy. Since the global synodic representation provides limited insight into the relative behavior of the trajectories, the analysis is mainly based on the time histories of the state components. A more direct geometric interpretation is later provided through the LVLH relative-motion plots.

As described in Chapter 6, the natural reference is updated online from the active stable manifold candidate and supplied to the artificial-reference NMPC problem. In this section, the resulting closed-loop evolution is analyzed through the time histories of the real trajectory X, the artificial trajectory X_s, and the natural reference $\hat{X}$.

Figures 7.1 and 7.2 show the time evolution of the position and velocity components. The plots compare the real spacecraft trajectory, the artificial trajectory, and the natural reference over the full simulation interval. The vertical black dashed lines indicate accepted switching events between manifold trajectories.

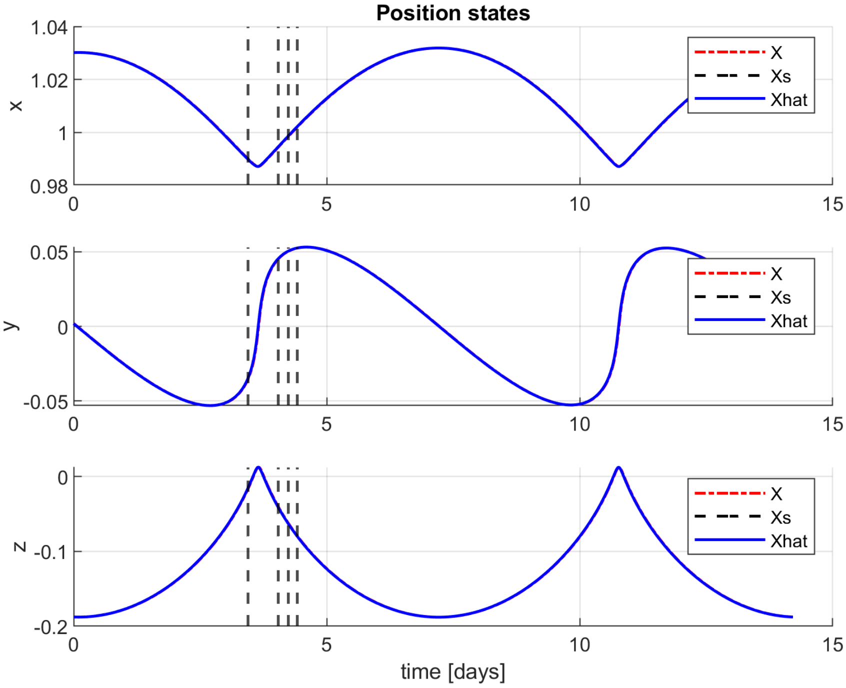

Figure 7.1: Position components of the real trajectory X, artificial trajectory X_s, and natural reference $\hat{X}$ during the closed-loop simulation.

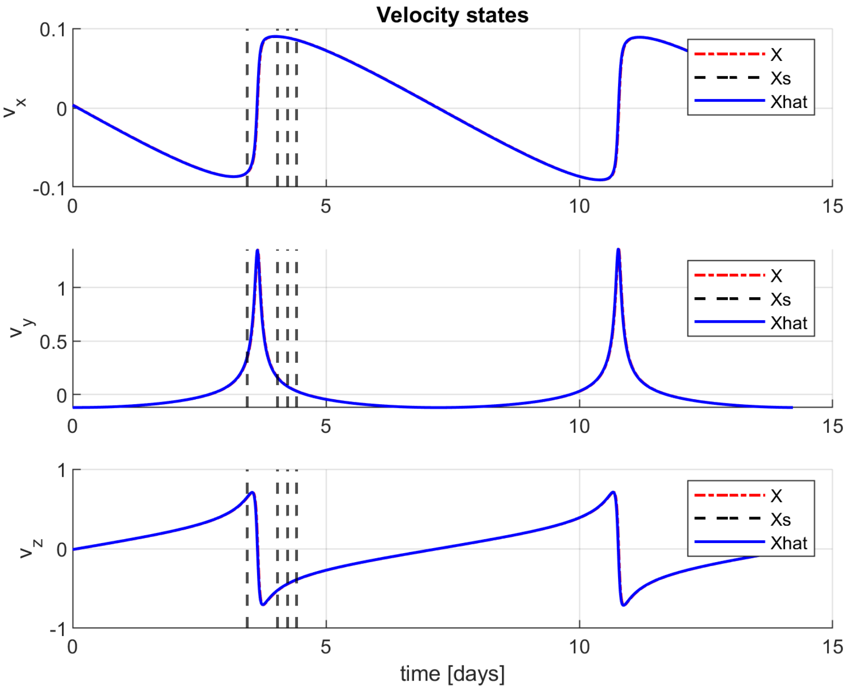

Figure 7.2: Velocity components of the real trajectory X, artificial trajectory X_s, and natural reference $\hat{X}$ during the closed-loop simulation.

The position and velocity histories show that the real and artificial trajectories remain closely related to the selected natural reference throughout the simulation. The velocity components are especially useful for verifying that the spacecraft is not only close to the reference in space, but also dynamically aligned with its motion.

A more detailed assessment of the relative tracking performance is provided in the following section through the corresponding position and velocity errors.

7.6 Tracking Errors

This section analyzes the tracking errors between the real spacecraft trajectory X, the artificial trajectory X_s, and the natural reference $\hat{X}$. While the previous section compared the individual state components, the present analysis focuses on the relative position and velocity errors.

The position errors are defined as

$$e^r_{X,X_s} = \|\mathbf{r} - \mathbf{r}_s\| , \qquad e^r_{X_s,\hat{X}} = \|\mathbf{r}_s - \hat{\mathbf{r}}\| , \qquad e^r_{X,\hat{X}} = \|\mathbf{r} - \hat{\mathbf{r}}\| . \tag{7.10}$$

Similarly, the velocity errors are

$$e^v_{X,X_s} = \|\mathbf{v} - \mathbf{v}_s\| , \qquad e^v_{X_s,\hat{X}} = \|\mathbf{v}_s - \hat{\mathbf{v}}\| , \qquad e^v_{X,\hat{X}} = \|\mathbf{v} - \hat{\mathbf{v}}\| . \tag{7.11}$$

Figure 7.3 shows the position tracking errors over the complete simulation interval. The first error, e^r_{X,X_s}, measures how closely the real spacecraft follows the optimized artificial trajectory. The second error, $e^r_{X_s,\hat{X}}$, quantifies the separation between the artificial trajectory and the selected natural reference. The third error, $e^r_{X,\hat{X}}$, provides the direct position error between the real trajectory and the natural reference.

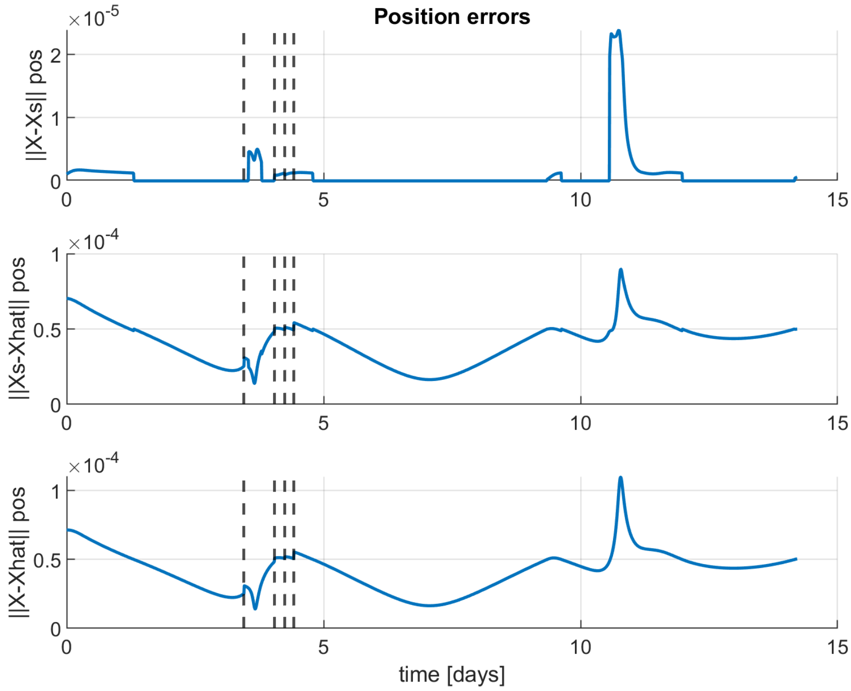

Figure 7.3: Position tracking errors between the real trajectory X, the artificial trajectory X_s, and the natural reference $\hat{X}$.

The position-error histories show that the real and artificial trajectories remain close to each other for most of the simulation. The error e^r_{X,X_s} stays at a lower level than the direct error with respect to the natural reference, confirming that the artificial trajectory acts as an intermediate path rather than forcing the real spacecraft to track $\hat{X}$ directly.

Local variations are visible near the switching instants, where the selected manifold candidate changes. However, the largest transient in the position errors occurs around $t \simeq 11$ days and is associated with the perilune passage. In this region, the spacecraft motion is faster and the natural dynamics are more sensitive, so a small mismatch in the reference-following process produces a larger temporary error.

Figure 7.4 reports the corresponding velocity errors. These quantities are particularly relevant because proximity in position alone is not sufficient to ensure that the spacecraft follows the same natural motion as the selected manifold trajectory.

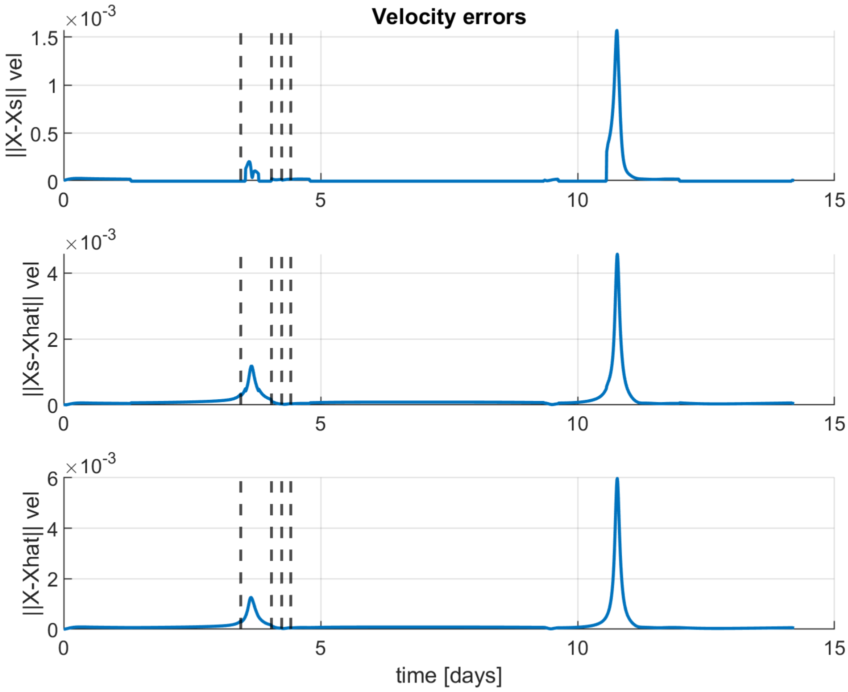

Figure 7.4: Velocity tracking errors between the real trajectory X, the artificial trajectory X_s, and the natural reference $\hat{X}$.

The velocity-error plot exhibits the same general behavior observed in the position errors. The real-to-artificial error remains small for most of the simulation, whereas the errors involving the natural reference are more sensitive to changes in the selected manifold segment and to the local dynamical conditions. The largest velocity transient appears around $t \simeq 11$ days, consistently with the perilune passage and with the corresponding peak observed in the position-error plot.

Overall, the tracking-error analysis confirms the role of the artificial trajectory as a smoothing layer between the real spacecraft motion and the natural manifold reference. The real trajectory remains closely connected to X_s, while X_s absorbs part of the mismatch with respect to $\hat{X}$. The most pronounced error peaks occur in dynamically demanding regions of the trajectory, rather than being caused only by the switching logic.

7.7 Control Effort and Accumulated Delta-v

This section analyzes the control effort required by the proposed manifold-based artificial-reference NMPC strategy. The control histories provide insight into when active correction maneuvers are required and how the coasting logic affects the overall propellant consumption.

Figure 7.5 shows the three components of the real control input U and of the artificial control input U_s. The horizontal dashed lines indicate the imposed component-wise bounds. For the real spacecraft input, the admissible interval is

$$-u_{\max} \leq u_i \leq u_{\max}, \qquad u_{\max} = 2 \times 10^{-3}, \tag{7.12}$$

while the artificial input is bounded by

$$-u_{s,\max} \leq u_{s,i} \leq u_{s,\max}, \qquad u_{s,\max} = 1.8 \times 10^{-3}. \tag{7.13}$$

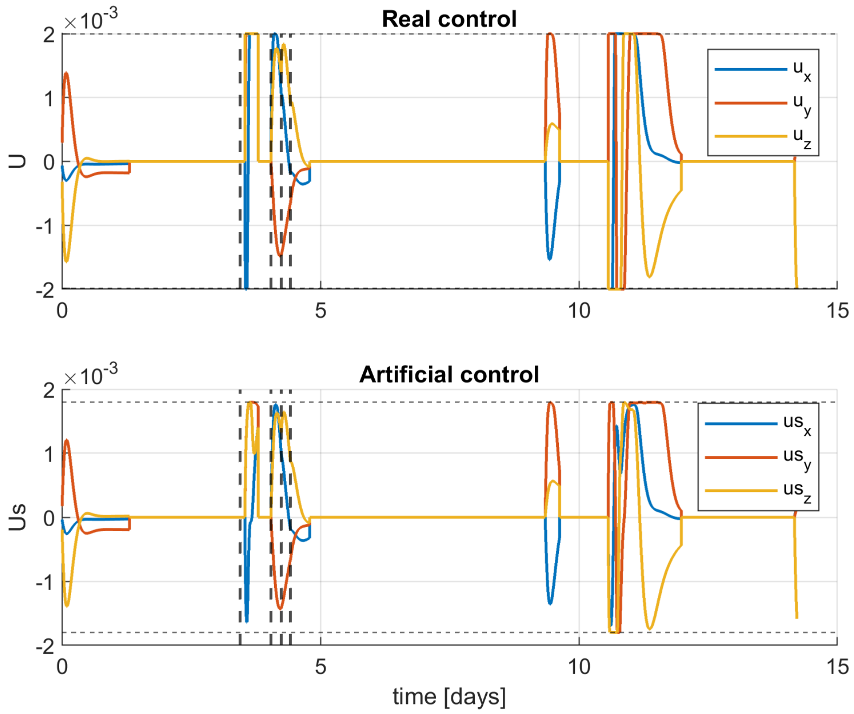

Figure 7.5: Real and artificial control components during the closed-loop simulation. The horizontal dashed lines indicate the imposed input bounds.

The control histories show that the input is not continuously active over the full simulation. Instead, most of the control effort is concentrated in localized intervals. An initial correction is required to compensate for the prescribed perturbation with respect to the selected manifold trajectory. Additional control peaks occur near the perilune passages, where the spacecraft motion is faster and the natural dynamics are more sensitive. In these intervals, some control components reach their admissible bounds, indicating that the input constraints are active.

Between these correction phases, the control remains close to zero for extended time intervals. This behavior is consistent with the role of the stable manifold reference and of the coasting logic: when the spacecraft motion remains sufficiently compatible with the natural dynamics, no corrective acceleration is required.

Figure 7.6 shows the Euclidean norms of the real and artificial control inputs, together with the corresponding running Δv estimates. The dimensional Δv is computed by converting the nondimensional acceleration input into physical units and integrating its norm over the sampling intervals:

$$\Delta v(t_\ell) = \sum_{i=0}^{\ell-1} \|\mathbf{u}_i\|\, a_{\text{unit}} \Delta t_{\text{dim}}, \tag{7.14}$$

where

$$a_{\text{unit}} = \frac{DU}{TU^2}, \qquad \Delta t_{\text{dim}} = T_s TU. \tag{7.15}$$

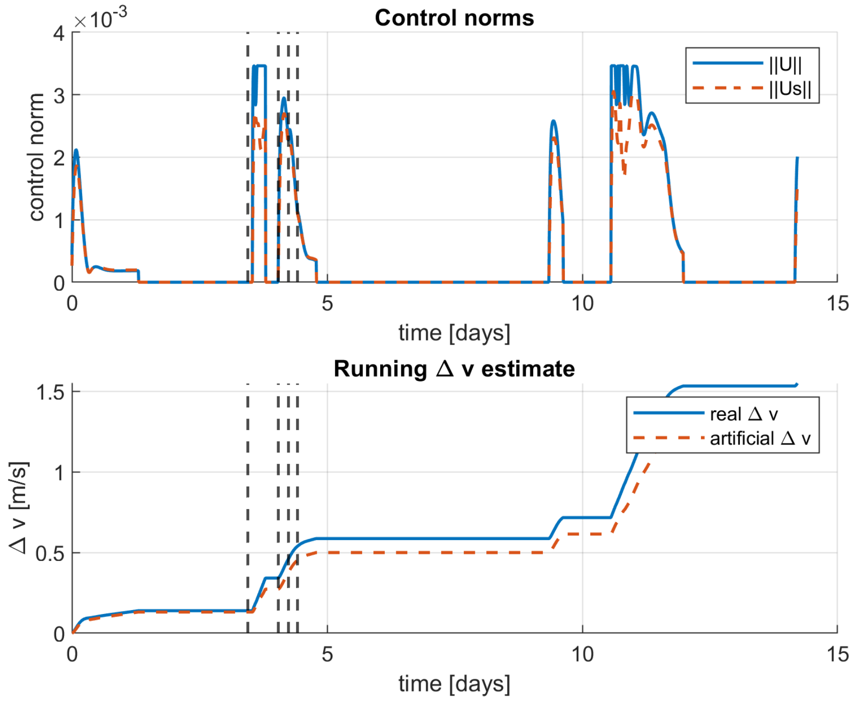

Figure 7.6: Control norms and accumulated Δv estimate for the real and artificial inputs.

The control-norm plot confirms that the main actuation demands occur in short time intervals. The most significant peaks are associated with perilune passages, where the spacecraft experiences faster motion and increased sensitivity to small deviations. These peaks are reflected in the accumulated Δv curve, which increases mainly in discrete steps rather than continuously.

The running Δv remains almost constant during long intervals in which the control norm is close to zero. This indicates that the spacecraft is able to exploit the natural CR3BP dynamics for significant portions of the simulation. The final accumulated real Δv is approximately

$$\Delta v_{\text{real}} \simeq 1.5 \text{ m/s}. \tag{7.16}$$

The artificial Δv is not a physical propellant cost, since U_s is associated with the optimized artificial trajectory rather than with the real spacecraft. However, it is useful as a diagnostic quantity because it indicates how much control authority is required to shape the artificial reference. Its evolution remains close to the real Δv trend, confirming that the artificial trajectory evolves consistently with the control effort applied to the real system.

Overall, the control analysis shows that the proposed strategy concentrates the required actuation in limited correction phases, while allowing the spacecraft to coast for extended portions of the trajectory. This behavior is consistent with the objective of exploiting stable manifold dynamics rather than continuously forcing the spacecraft to track an externally prescribed reference.

7.8 Jacobi Constant Evolution

This section analyzes the evolution of the Jacobi constant along the real trajectory X, the artificial trajectory X_s, and the natural reference $\hat{X}$. Since no Jacobi-related terms are included in the cost function adopted for these simulations, the Jacobi constant is not directly penalized within the NMPC problem. It is therefore used here as a diagnostic quantity to evaluate the energy-like behavior of the controlled and reference trajectories.

In the uncontrolled CR3BP, the Jacobi constant is an integral of motion. As a consequence, a periodic orbit and its associated invariant manifolds belong to the same Jacobi level. Since the natural reference $\hat{X}$ is generated from the stable manifold structure and propagated under uncontrolled CR3BP dynamics, its Jacobi constant remains essentially constant over time, apart from numerical integration errors.

Figure 7.7 shows the values of $C(X)$, $C(X_s)$, and $C(\hat{X})$ over the closed-loop simulation.

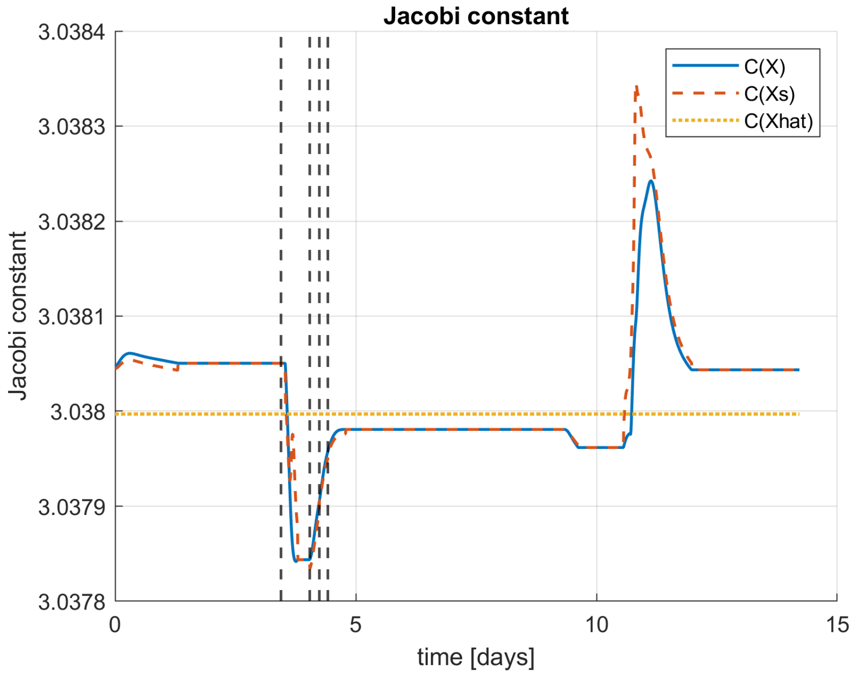

Figure 7.7: Jacobi constant evaluated along the real trajectory X, artificial trajectory X_s, and natural reference $\hat{X}$.

The Jacobi constant of the natural reference remains nearly constant throughout the simulation. This is consistent with the fact that the reference belongs to the same energy level as the underlying halo orbit and its associated stable manifold.

By contrast, the real and artificial trajectories exhibit variations in their Jacobi constant. This behavior is expected because both X and X_s are governed by controlled dynamics. The Jacobi integral is conserved only in the absence of control, whereas the applied acceleration modifies the energy-like level associated with the trajectory.

The curves of $C(X)$ and $C(X_s)$ remain close to each other for most of the simulation, consistently with the role of the artificial trajectory as an intermediate reference for the real spacecraft. More pronounced variations occur during the main correction phases, especially near dynamically demanding regions of the trajectory such as perilune passages. These variations are consistent with the peaks observed in the control histories.

Overall, the Jacobi-constant evolution confirms that the natural reference preserves the Jacobi level of the underlying halo orbit and stable manifold, while the controlled real and artificial trajectories undergo limited variations due to the applied control inputs. Since the Jacobi constant is not directly penalized in this configuration, these variations are monitored but not explicitly suppressed by the optimizer.

7.9 LVLH Relative Motion Analysis

This section analyzes the closed-loop motion in a Local Vertical Local Horizontal frame centered on the moving natural reference $\hat{X}$. This representation is useful because it removes the dominant motion of the reference trajectory and highlights the relative behavior of the real spacecraft trajectory X and of the artificial trajectory X_s.

The LVLH frame is constructed at each time instant using the position and velocity of the natural reference with respect to the Moon. The origin of the frame is placed at $\hat{X}$, while the relative positions of X and X_s are expressed in the local V-bar, R-bar, and H-bar directions. Therefore, the point

$$\mathbf{0} = \begin{bmatrix} 0 & 0 & 0 \end{bmatrix}^T \tag{7.17}$$

represents the instantaneous position of the moving reference $\hat{X}$.

Figure 7.8 shows the three-dimensional relative motion of $X - \hat{X}$ and $X_s - \hat{X}$. The artificial trajectory remains close to the real trajectory for most of the simulation, confirming that X_s provides an intermediate path between the controlled spacecraft and the natural reference. The relative displacement remains within a range of a few tens of kilometers, which is more clearly visible in the LVLH frame than in the global synodic representation.

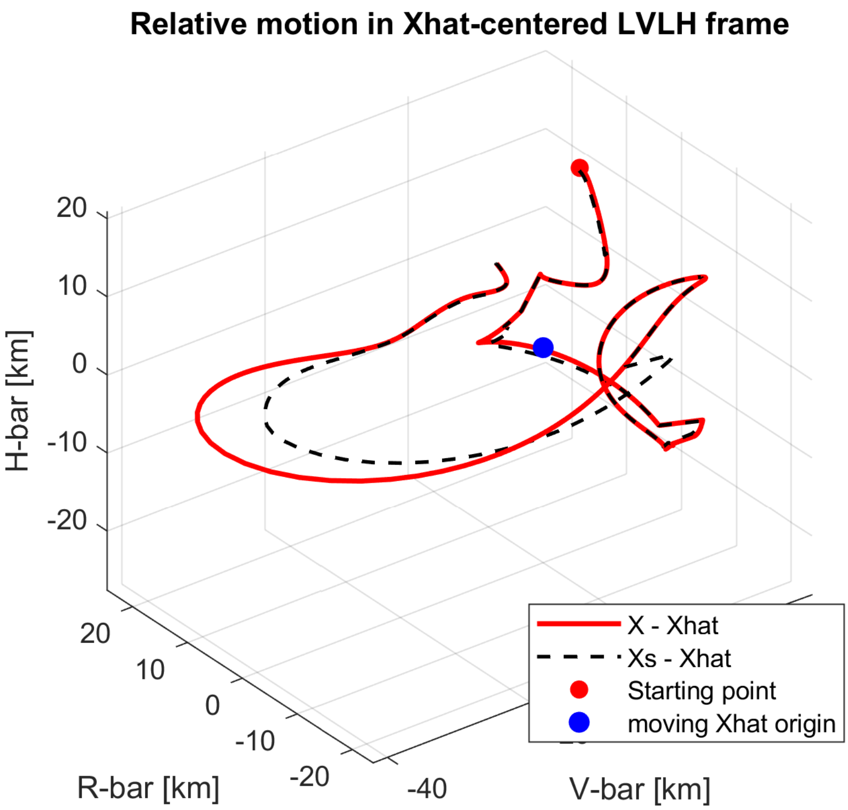

Figure 7.8: Relative motion of the real trajectory X and artificial trajectory X_s with respect to the moving natural reference $\hat{X}$ in the LVLH frame.

Figures 7.9 and 7.10 show the corresponding planar projections. The R-bar/V-bar projection highlights the in-plane relative motion, while the H-bar/V-bar projection shows the out-of-plane relative behavior.

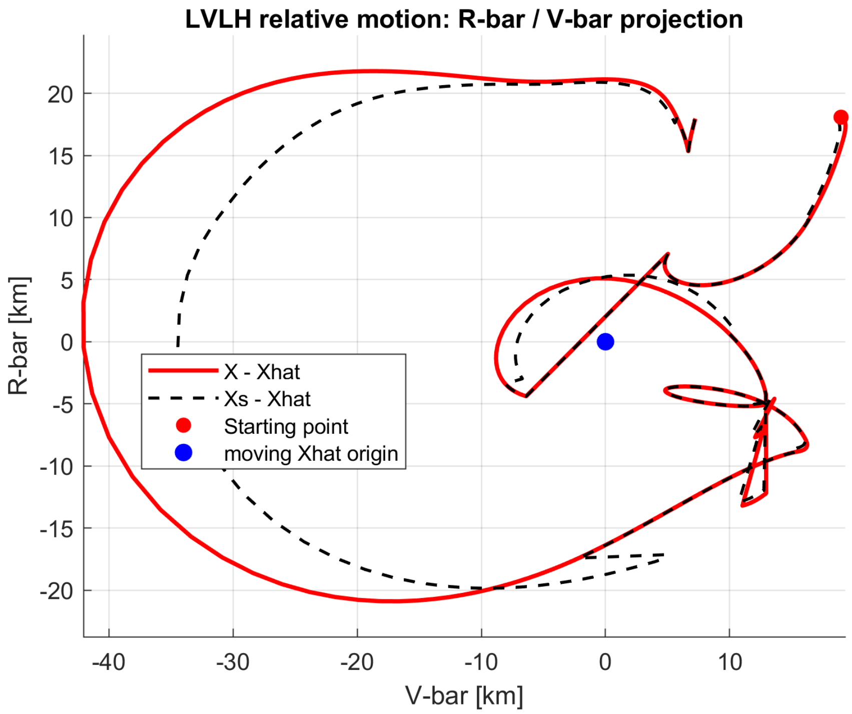

Figure 7.9: LVLH relative motion projected onto the R-bar/V-bar plane.

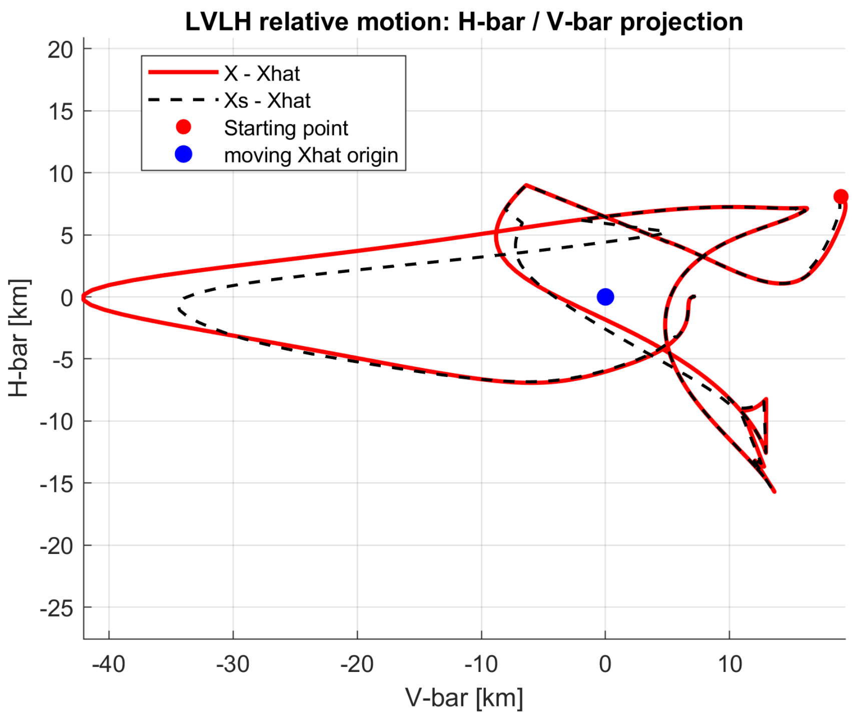

Figure 7.10: LVLH relative motion projected onto the H-bar/V-bar plane.

The planar projections show that the real and artificial relative trajectories have similar overall shapes, although local separations appear in some portions of the motion. This behavior is consistent with the artificial-reference formulation: the artificial trajectory is not forced to coincide exactly with the real spacecraft trajectory, but it remains sufficiently close to provide a dynamically feasible intermediate reference.

The largest relative excursions occur in regions where the dynamics are more demanding, consistently with the peaks observed in the tracking errors and in the control histories. In particular, the LVLH plots provide a clearer geometric interpretation of the transient relative motion that is not easily visible in the synodic-frame representation.

Overall, the LVLH analysis confirms that the proposed controller maintains the real spacecraft trajectory in the neighborhood of the moving natural reference, while allowing the artificial trajectory to act as a smoothing layer between the real motion and the selected manifold-based reference.

7.10 Comparison with Conventional NMPC

This section compares the proposed manifold-based artificial-reference NMPC strategy with a conventional NMPC formulation. In the conventional case, the controller directly tracks the target halo orbit, without exploiting the stable manifold library and without introducing an artificial trajectory.

The same CR3BP model, discretization scheme, prediction horizon, input bounds, and initial perturbed condition are adopted. The manifold is used only to define the initial condition, so that both controllers start from the same perturbed state. This allows the effect of the reference architecture to be assessed more clearly.

In the proposed strategy, the reference is generated from the stable manifold structure and the real trajectory is connected to it through the artificial trajectory. In the conventional formulation, instead, the reference is selected directly on the halo orbit as the nearest halo point to the current spacecraft state.

Tracking performance

Figure 7.11 shows the position and velocity errors obtained with the conventional NMPC formulation. The errors are evaluated with respect to the nearest point on the target halo orbit, denoted by X_{ref}.

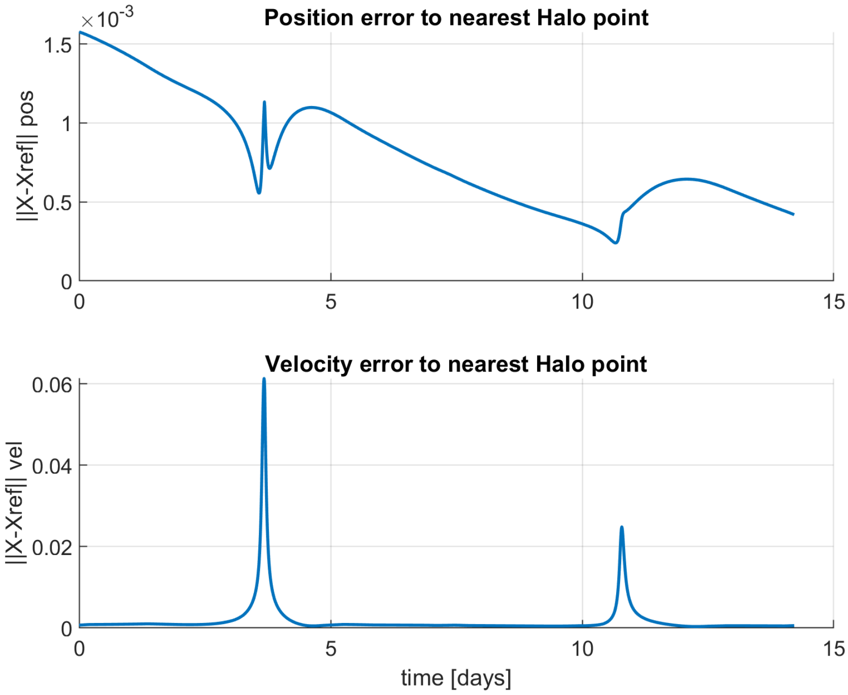

Figure 7.11: Position and velocity errors obtained with the conventional NMPC formulation, evaluated with respect to the nearest point on the target halo orbit.

The position error decreases during the simulation, showing that the conventional controller is able to bring the spacecraft closer to the halo orbit. However, the error remains affected by oscillations and does not vanish monotonically. The velocity error exhibits more pronounced transients, especially near dynamically demanding portions of the trajectory.

Compared with the manifold-based artificial-reference strategy, the conventional controller shows a more aggressive tracking behavior. This is expected because the spacecraft is initialized close to a stable manifold arc rather than close to the halo orbit itself. Therefore, direct tracking of the halo orbit requires the controller to compensate for a larger dynamical mismatch from the beginning of the simulation.

Control effort and accumulated Δv

The difference between the two strategies becomes more evident when the control effort is analyzed. Figure 7.12 shows the real control components of the conventional NMPC formulation, together with the corresponding dimensional acceleration components.

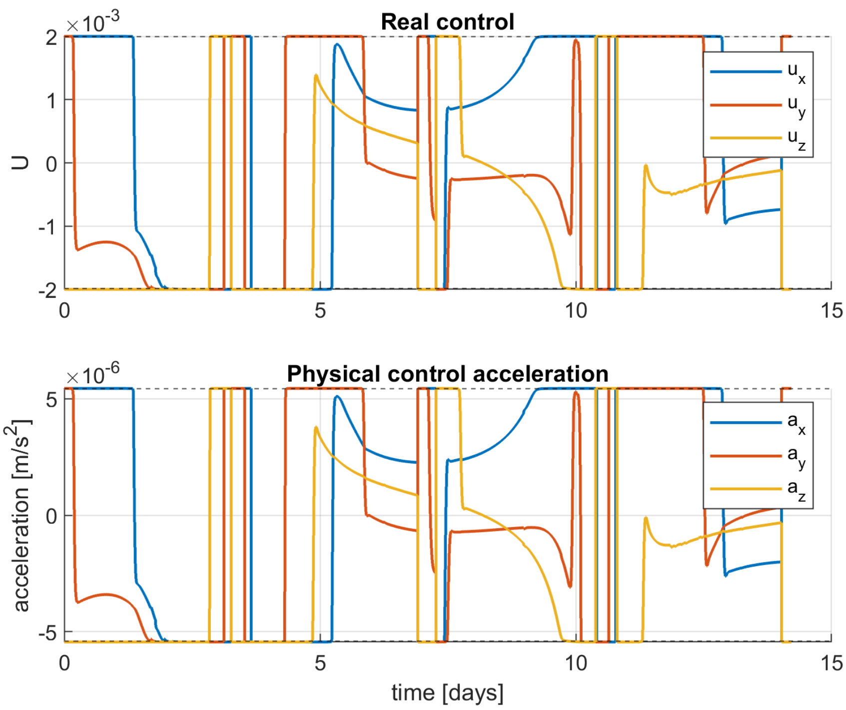

Figure 7.12: Control components and corresponding dimensional accelerations for the conventional NMPC formulation.

The conventional controller frequently drives the input components to their admissible bounds. Unlike the proposed manifold-based strategy, where control is concentrated in localized correction phases, the conventional formulation requires sustained actuation over a large portion of the simulation. This is a direct consequence of forcing the spacecraft to track the halo orbit without first exploiting the stable manifold dynamics.

Figure 7.13 reports the norm of the control input and the corresponding accumulated Δv.

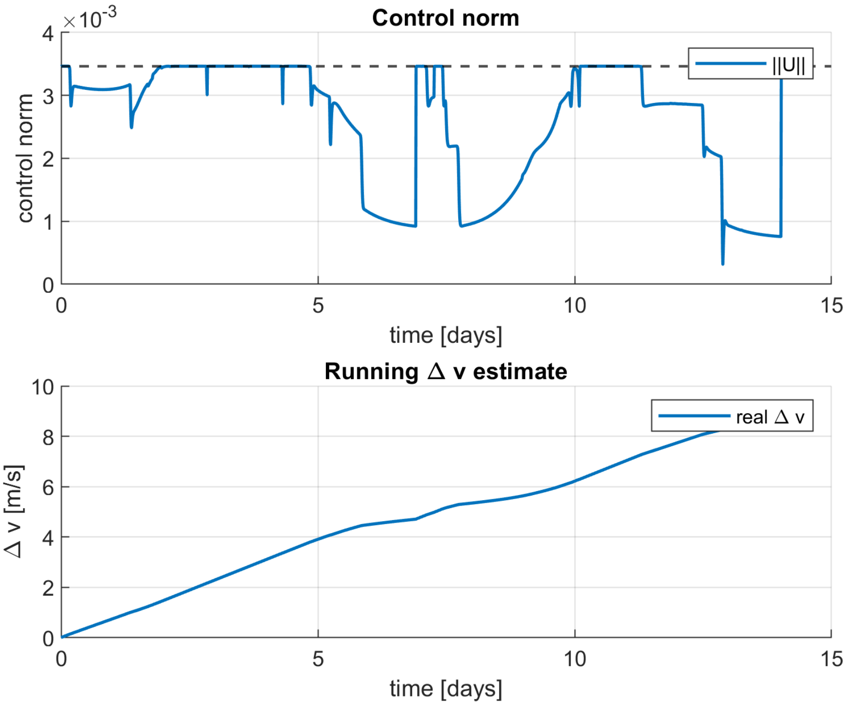

Figure 7.13: Control norm and accumulated Δv estimate for the conventional NMPC formulation.

The control norm remains high for most of the simulation and often approaches the maximum admissible value. As a result, the accumulated Δv grows almost continuously. The final value is approximately

$$\Delta v_{\mathrm{conv}} \simeq 8.5\text{–}9 \text{ m/s}. \tag{7.18}$$

This value is substantially larger than the one obtained with the proposed strategy, for which the final real Δv is approximately

$$\Delta v_{\mathrm{prop}} \simeq 1.5 \text{ m/s}. \tag{7.19}$$

The comparison therefore shows that the use of the stable manifold reference and of the artificial trajectory significantly reduces the required control effort. In the proposed formulation, the spacecraft is not forced to track the halo orbit directly from the perturbed initial condition; instead, it is guided along a dynamically meaningful manifold-based path.

Relative motion in the LVLH frame

The LVLH representation provides a geometric comparison of the relative motion produced by the two controllers. For each controller, the LVLH frame is centered on the moving reference used by that controller. In the proposed strategy, this reference is the selected natural manifold reference $\hat{X}$. In the conventional NMPC formulation, the reference is the nearest point on the target halo orbit, denoted by X_{ref}.

Figures 7.14, 7.15, and 7.16 show the relative motion obtained with the conventional NMPC formulation. The trajectory is represented as $X - X_{\mathrm{ref}}$ in the V-bar, R-bar, and H-bar directions.

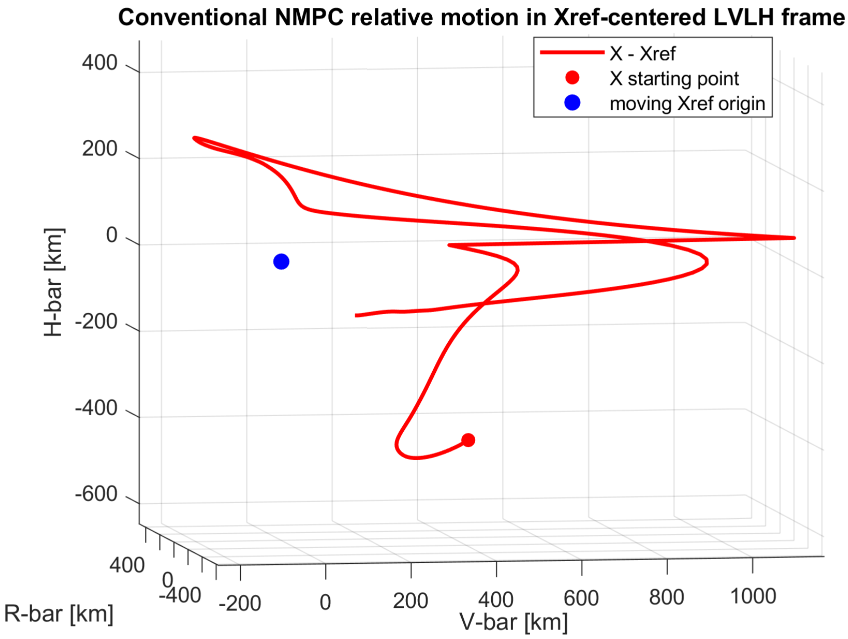

Figure 7.14: Relative motion of the conventional NMPC trajectory with respect to the moving halo reference X_{ref} in the LVLH frame.

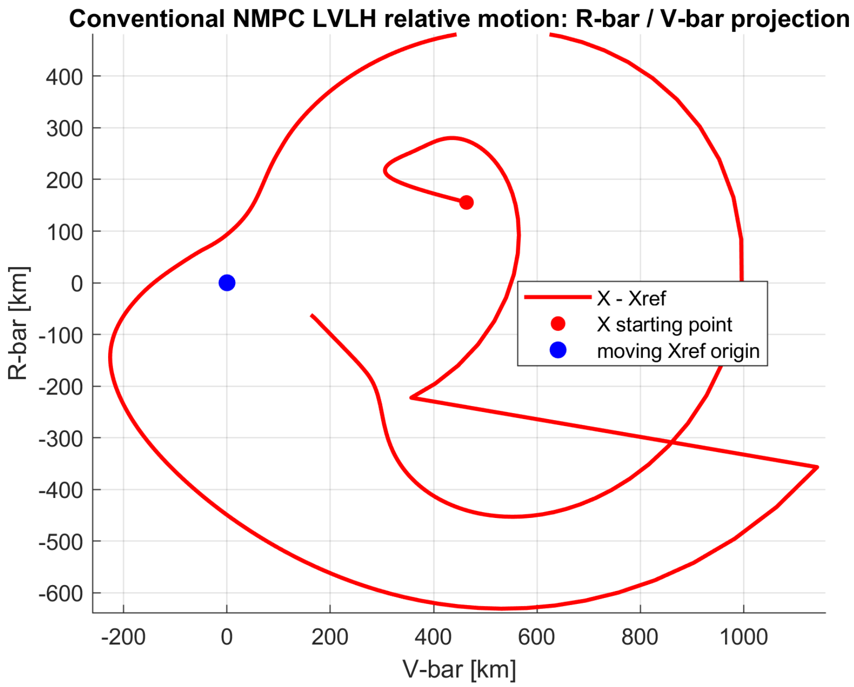

Figure 7.15: LVLH relative motion of the conventional NMPC trajectory projected onto the R-bar/V-bar plane.

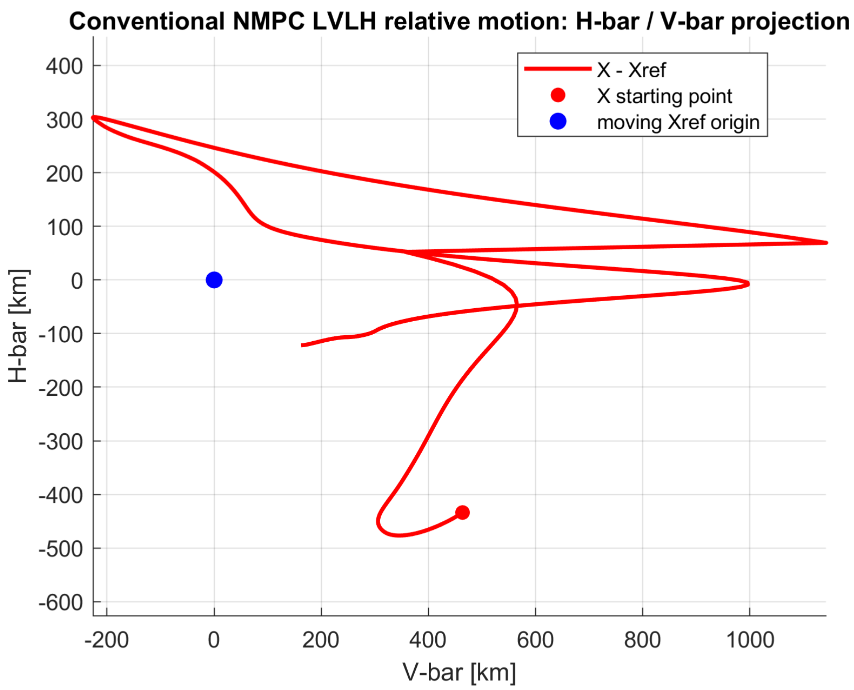

Figure 7.16: LVLH relative motion of the conventional NMPC trajectory projected onto the H-bar/V-bar plane.

The relative excursions in the conventional case are much larger than those observed with the proposed strategy. In the conventional NMPC case, the relative motion spans several hundreds of kilometers in the LVLH frame. By contrast, the manifold-based artificial-reference strategy keeps the relative motion confined to a much smaller region around the moving natural reference.

This difference confirms that direct halo tracking is not well aligned with the considered initial condition. Since the spacecraft starts close to a stable manifold arc, forcing it to track the halo orbit directly produces large relative excursions and requires stronger corrective actions. The proposed method instead exploits the stable manifold as an intermediate dynamical structure, which naturally connects the initial region to the neighborhood of the halo orbit.

The comparison shows that direct halo tracking leads to a more demanding control problem, whereas the proposed strategy benefits from the use of the stable manifold as an intermediate natural path. This results in lower control effort, smaller relative excursions, and a reduced accumulated Δv.

7.11 Discussion

The numerical results show a clear difference between the proposed manifold-based artificial-reference NMPC strategy and the conventional NMPC formulation. In the proposed case, the real trajectory remains closely connected to the artificial trajectory, while the artificial trajectory follows the selected manifold-based natural reference. This behavior is visible in the position and velocity histories, where X, X_s, and $\hat{X}$ remain close for most of the simulation.

The tracking-error plots confirm the role of the artificial trajectory as an intermediate layer. The real-to-artificial errors remain smaller than the direct errors with respect to the natural reference, showing that the controller does not force the spacecraft to track the manifold reference rigidly. The largest transients occur near dynamically sensitive regions, especially close to perilune passages, where both the tracking errors and the control effort increase.

The control plots highlight the main advantage of the proposed strategy. The input is concentrated in localized correction phases, with long intervals of nearly zero control. This leads to a final accumulated Δv of approximately 1.5 m/s. By contrast, the conventional NMPC controller frequently drives the input to its bounds and requires sustained actuation over a large portion of the simulation, resulting in an accumulated Δv of approximately 8.5–9 m/s.

The LVLH representations provide a geometric confirmation of this difference. In the proposed strategy, the relative motion with respect to the moving natural reference remains confined to a region of a few tens of kilometers. In the conventional formulation, the relative motion with respect to the moving halo reference spans several hundreds of kilometers. This indicates that direct halo tracking is less aligned with the considered initial condition, which is closer to a stable manifold arc than to the halo orbit itself.

Table 7.4: Qualitative comparison between the proposed manifold-based artificial-reference NMPC strategy and the conventional NMPC formulation.

Metric	Proposed strategy	Conventional NMPC
Reference type	Stable manifold reference $\hat{X}$	Nearest halo point X_{ref}
Tracking structure	$X \to X_s \to \hat{X}$	$X \to X_{\mathrm{ref}}$
Artificial trajectory	Included	Not included
Stable manifold exploitation	Yes	No
Control behavior	Localized correction phases	Persistently active
Input saturation	Mainly during correction phases	Frequent over the trajectory
Final accumulated Δv	$\simeq 1.5$ m/s	$\simeq 8.5$–9 m/s
LVLH relative motion	Few tens of kilometers	Several hundreds of kilometers
Main limitation	Transients near sensitive regions	High control demand

Overall, the proposed strategy exploits the stable manifold as a dynamically meaningful intermediate path and uses the artificial trajectory to avoid aggressive direct tracking. This results in smaller relative excursions, less continuous actuation, and a significantly lower accumulated Δv than the conventional NMPC formulation.

7.12 Concluding Remarks

This chapter presented the numerical validation of the proposed manifold-based artificial-reference NMPC strategy in the Earth–Moon CR3BP. The results showed that the controller can guide the spacecraft toward the target halo orbit while exploiting the natural stable manifold structure.

The proposed strategy reduced the need for continuous actuation and achieved a lower accumulated Δv than the conventional NMPC formulation considered for comparison. The following chapter summarizes the main conclusions of the thesis and outlines possible directions for future work.

Chapter 8

Conclusions and Future Work

8.1 Conclusions

This thesis investigated a manifold-based nonlinear model predictive control strategy for spacecraft guidance in the Earth–Moon Circular Restricted Three-Body Problem. The main objective was to exploit the natural transport structures associated with stable invariant manifolds of halo orbits, while retaining the flexibility of a constrained predictive control formulation.

The dynamical model was first introduced in the nondimensional Earth–Moon CR3BP, using the synodic reference frame and a controlled first-order state-space formulation. The Jacobi constant was discussed as a fundamental invariant of the uncontrolled dynamics, and a fourth-order Runge–Kutta discretization was adopted to obtain the discrete-time model used in the controller and in the manifold propagation.

The invariant structures associated with libration-point dynamics were then analyzed, with particular attention to halo orbits and their stable manifolds. A stable manifold library was generated offline and used as the basis for the online reference-generation strategy. This allowed the controller to exploit trajectories that are naturally connected to the target periodic orbit.

The proposed guidance strategy combined the manifold-based reference structure with an artificial-reference NMPC formulation. Instead of forcing the real spacecraft trajectory to directly track a prescribed natural reference, an optimized artificial trajectory was introduced as an intermediate path. This provided an additional degree of freedom in the optimization problem and allowed a more flexible connection between the real spacecraft motion and the selected natural reference.

The numerical simulations showed that the proposed strategy can guide the spacecraft from a perturbed initial condition toward the target halo orbit while exploiting the stable manifold dynamics. The control effort was mainly concentrated in localized correction phases, while extended portions of the trajectory were characterized by nearly zero input. The accumulated Δv was significantly lower than that obtained with a conventional NMPC controller directly tracking the halo orbit.

The comparison with the conventional formulation highlighted the advantage of using a reference architecture aligned with the geometry of the problem. Since the initial condition was located close to a stable manifold arc rather than close to the halo orbit itself, direct halo tracking required sustained control activity and larger relative excursions. By contrast, the manifold-based artificial-reference strategy reduced the need for aggressive control action by using the stable manifold as a natural intermediate path.

Overall, the results support the effectiveness of combining invariant-manifold structures with artificial-reference NMPC for spacecraft guidance in multi-body dynamical environments.

8.2 Future Work

Several extensions can be considered to further develop the proposed approach. A first direction concerns the enlargement of the stable manifold library. The present implementation considered a limited number of candidate trajectories, while a denser library including different branches, halo orbit samples, and energy levels could improve the flexibility of the online reference-selection process.

A further development in this direction would be the transition from an offline precomputed library to an online manifold-generation strategy. In the present formulation, the stable manifold trajectories are generated before the closed-loop simulation and then used as fixed reference candidates. An online approach would instead generate or update manifold trajectories during the guidance process, based on the current spacecraft state, the selected target orbit, or the evolution of the optimization problem. This could reduce the dependence on a fixed database and allow the controller to adapt the available natural references to changing mission conditions. However, such an approach would also increase the computational burden, since manifold generation requires the computation of local eigendirections and the numerical propagation of the associated perturbed states.

Another possible extension concerns the formulation of the guidance problem in a modal coordinate representation rather than directly in the physical synodic state variables. In this thesis, reference selection and tracking are based on weighted errors in the Cartesian state space, involving position and velocity components. Although this representation is intuitive and directly related to the spacecraft motion, it does not explicitly distinguish between dynamically different error components.

A modal approach could instead project the relative state error onto local eigendirections associated with the reference trajectory, separating stable, unstable, and center-like components of the motion. In this way, the controller would not only evaluate the magnitude of the deviation from the reference, but also the dynamical nature of that deviation. For example, an error component along an unstable direction may require stronger correction than an error of comparable size along a naturally stable direction. Such information could be used to define more dynamically meaningful tracking costs, switching criteria, or terminal constraints.

The main limitation of a modal formulation is that the modal basis is local and depends on the reference trajectory. Therefore, the eigendirections would have to be computed and updated consistently along the orbit or manifold. Moreover, since the CR3BP dynamics are nonlinear, an eigendirection-based decomposition is accurate only locally. Large deviations from the reference may not be well represented by a linear modal approximation. Despite these challenges, a modal representation could improve the controller's ability to distinguish dynamically critical deviations from errors that are naturally reduced by the system dynamics.

A further extension concerns the inclusion of the arrival time as an optimization variable. In this thesis, the guidance strategy focuses on tracking stable manifold trajectories in order to reach the prescribed halo orbit, while the timing of the motion is mainly dictated by the selected reference and by the fixed prediction discretization. However, in some mission scenarios it may be important not only to reach the target orbit, but also to arrive at a specific point along it or at a prescribed orbital phase. Introducing the arrival time, or equivalently the phase along the target orbit, as an additional optimization variable could provide greater flexibility in the guidance problem. This would allow the controller to select both the path and the timing of the approach, improving its capability to satisfy rendezvous-like requirements or mission constraints associated with a specific location on the periodic orbit.

The switching logic could also be improved. In this thesis, trajectory switching is based on a weighted distance between the spacecraft state and the candidate manifold trajectories, together with a hold condition to avoid rapid switching. Future work could make this decision more predictive. For example, instead of selecting an alternative manifold only because it is closer to the current spacecraft state, the controller could also evaluate the expected NMPC cost, the required control effort, or the feasibility of the terminal constraint over the prediction horizon. In this way, the selected manifold would not necessarily be the closest one at the current instant, but the one expected to produce the best closed-loop behavior.

The coasting strategy could also be refined. The current implementation activates coasting through a one-step uncontrolled prediction and fixed position and velocity tolerances. Multi-step coasting checks or adaptive thresholds could improve the ability of the controller to exploit natural dynamics while maintaining accurate tracking performance.

Another important extension is the inclusion of model uncertainties and external perturbations. The present study is based on the ideal CR3BP model, whereas real mission scenarios require higher-fidelity dynamics, navigation errors, actuator limitations, and additional perturbing effects. Robust or stochastic NMPC formulations could be investigated to address these aspects.

Finally, the proposed strategy could be tested within a higher-fidelity mission design framework, including ephemeris dynamics and operational constraints. This would allow a more realistic assessment of the method and of its applicability to autonomous guidance near libration-point orbits.

Bibliography

[1] Nicola Sullo, Priscilla A. De Sousa Silva, Maisa O. Terra, and Matteo Ceriotti. Optimisation of low-thrust and hybrid earth–moon transfers. In *Proceedings of the 67th International Astronautical Congress*, Guadalajara, Mexico, 2016. IAC 2016, 26–30 September 2016.

[2] Wang Sang Koon, Martin W. Lo, Jerrold E. Marsden, and Shane D. Ross. *Dynamical Systems, the Three-Body Problem and Space Mission Design*. 2022.

[3] D. Limón, I. Alvarado, T. Alamo, and E. F. Camacho. MPC for tracking piecewise constant references for constrained linear systems. *Automatica*, 44(9):2382–2387, 2008.

[4] D. L. Richardson. Analytic construction of periodic orbits about the collinear points. *Celestial Mechanics*, 22(3):241–253, 1980.

[5] Kathleen C. Howell. Three-dimensional, periodic, "halo" orbits. *Celestial Mechanics*, 32:53–71, 1984.

[6] G. Gómez, A. Jorba, C. Simó, and A. Masdemont. Study of the transfer between halo orbits. *Acta Astronautica*, 43(9–10):493–520, 1998.

[7] James B. Rawlings, David Q. Mayne, and Moritz M. Diehl. *Model Predictive Control: Theory, Computation, and Design*. Nob Hill Publishing, Madison, WI, 2 edition, 2017.

[8] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789–814, 2000.

[9] Michele Pagone, Giordana Bucchioni, Francesco Alfino, and Carlo Novara. Autonomous lunar rendezvous trajectory planning and control using nonlinear MPC and pontryagin's principle. *IFAC Journal of Systems and Control*, 2024.

[10] Joel A. E. Andersson, Joris Gillis, Greg Horn, James B. Rawlings, and Moritz M. Diehl. CasADi: A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019.

[11] D. Limón, A. Ferramosca, I. Alvarado, and T. Alamo. Nonlinear MPC for tracking piece-wise constant reference signals. *IEEE Transactions on Automatic Control*, 63(11):3735–3750, 2018.

[12] Pablo Krupa, Johannes Köhler, Antonio Ferramosca, Ignacio Alvarado, Melanie N. Zeilinger, Teodoro Alamo, and Daniel Limón. Model predictive control for tracking using artificial references: Fundamentals, recent results and practical implementation, 2024.