



**Politecnico
di Torino**

Politecnico di Torino

Master's Degree in Aerospace Engineering

Hybrid Optimization Technique for Earth-Moon Transfer Trajectories Human and Cargo Missions to the Lunar Gateway

Supervisors

Prof. Lorenzo Casalino
Ing. Andrea D'ottavio
Dr. Giorgio Fasano

Candidate

Riccardo Roberto

Academic year 2024-2025

*Se tu proietti una grande
immagine di te, ti crei una
direzione, e vai in quella
direzione. Quindi è probabile che
in un modo o nell'altro ci
arriverai.*

JAGO

Acknowledgemnts

Al termine di questo percorso, desidero rivolgere un sentito ringraziamento a coloro che hanno contribuito alla realizzazione di questa tesi e alla mia crescita personale e accademica.

Un ringraziamento speciale va al mio relatore, il Professor Lorenzo Casalino, per la sua costante disponibilità, la sua guida preziosa e i suoi saggi consigli. La sua profonda conoscenza della materia e la sua passione per la ricerca sono state per me fonte di grande ispirazione.

Vorrei esprimere la mia più sincera gratitudine ad Andrea D'Ottavio. Che ha creduto fin da subito in questo progetto e mi ha guidato passo dopo passo nella sua completa realizzazione fornendomi preziosi consigli.

Un ringraziamento speciale a Giorgio Fasano, il cui prezioso aiuto e la cui profonda competenza tecnica sono stati indispensabili per il completamento di questo lavoro di tesi.

Infine, desidero ringraziare tutti i ragazzi dell'ufficio di Mission Analysis. L'ambiente stimolante e collaborativo che avete creato ha reso questo periodo un'esperienza indimenticabile e soprattutto piacevole. Vi ringrazio per gli infiniti momenti di confronto.

Abstract

This thesis addresses the issue of low-energy transfers between Low Earth Orbit (LEO) and Near Rectilinear Halo Orbit (NRHO) within the Earth–Moon system. The focus is on supporting the Lunar Gateway within NASA’s Artemis program. The cis-lunar domain is presented as an ideal environment for testing technologies, establishing logistics hubs, and enabling future Mars missions.

Trajectory optimization is constrained by the Kourou launch site in French Guiana as the departure point, which has strict altitude and inclination constraints. These limitations significantly impact the feasibility of different trajectories.

The optimization process employs a combination of two approaches. First, a Genetic Algorithm (GA) scans the solution space to identify promising trajectories. Then, a sequential quadratic programming (SQP) method uses the SNOPT solver to fine-tune these trajectories.

Since the cis-lunar area is highly unstable, the trajectories reach a specific point along the NRHO. The European Space Operations Centre (ESOC) and NASA’s Spacecraft, Planet, Instrument, C-matrix, Events (SPICE) tool provide precise ephemerides to calculate this point.

Simplified models were also tested to gain a panoramic view of the problem. In particular, we employed the Circular Restricted Three-Body Problem (CR3BP) with a single-shooting method. While not as precise as the high-fidelity n-body model, the CR3BP is useful for gaining an initial understanding of how the solution space behaves. It also helps identify primary sensitivities before conducting more complex simulations.

Since dealing with altitude and inclination constraints simultaneously can be challenging, a multiple-shooting setup was used. The trajectory was split into nodes, each with a discrete impulse and flexible segment duration. This setup significantly improves the solutions and makes the trajectory more adjustable. Additionally, the GA+SNOPT hybrid optimization technique was applied to refine the solution.

The work was later expanded by varying the departure date over the course of a year. This allowed for the identification of better launch windows along the NRHO. Some adjustments were also made to avoid clustering impulses too closely together, such as modifying the bounds of the variables governing the duration of individual segments. These improvements reduced the number of maneuvers required, allowing for lower fuel consumption. Other types of trajectories, such as

Weak Stability Boundary (WSB) trajectories, were employed to identify routes with minimal delta-V.

Overall, the results demonstrate that combining the two optimization methods effectively identifies low-cost trajectories.

Contents

Abstract	4
List of Figures	10
Acronyms	11
Introduction	12
1 Mission context	15
1.1 Access to cislunar space and the lunar Gateway	15
1.2 The low-energy transfer problem	16
2 Theoretical background and state of the art	17
2.1 Fundamentals of astrodynamics	17
2.1.1 The two-body problem	17
2.1.2 The restricted three-body problem (CR3BP)	18
2.1.3 Lagrange points and periodic orbits	19
2.1.4 Near-Rectilinear Halo Orbits (NRHO)	20
2.1.5 Perturbations and N-Body models	21
2.1.6 Role in mission design	21
2.2 Trajectory optimization methods	22
2.2.1 Direct vs. Indirect Approaches	22
2.2.2 Shooting techniques: single vs. multiple shooting	22
2.3 Optimization algorithms	24
2.3.1 Global optimization: genetic algorithms (GA)	25
2.3.2 Local optimization: sequential quadratic programming (SQP) and the SNOPT solver	25
2.4 Forward vs. backward propagation approaches	25
3 Methodology and problem formulation	28
3.1 High-fidelity dynamical model	28
3.1.1 N-Body propagation and numerical integration	29
3.1.2 Use of SPICE ephemerides (JPL/ESOC)	29

3.1.3	Reference frames and usage	30
3.2	Formulation of the optimization problem	33
3.2.1	Decision vector (Delta-V and flight times τ)	34
3.2.2	Objective function: minimization of total Delta-V	34
3.2.3	Boundary and path constraints (altitude, inclination, safety)	35
3.3	Hybrid optimization strategy (GA + SQP)	35
3.3.1	Phase 1: global search with genetic algorithm and multiple shooting	36
3.3.2	Phase 2: local refinement	36
3.4	Advanced strategies for solution space exploration	37
3.4.1	Parallel Multi-Start with stochastic perturbations	37
3.4.2	Path relinking technique for alternative local minima identification	37
3.4.3	Exploration of departure dates and optimal launch windows	38
3.4.4	Application of weak stability boundary (WSB) transfers	39
4	Implementation, Results, and Analysis	41
4.1	Tool overview and logic	41
4.1.1	Conventions (Frames, Time, Units, Metrics)	42
4.2	Methods shared by both cases	43
4.2.1	Problem formulation	43
4.2.2	GA+SNOPT pipeline	43
4.2.3	Diagnostics and outputs	43
4.3	Direct transfer case	45
4.3.1	GA results	45
4.3.2	SNOPT refinement	51
4.4	Cargo Case	55
4.4.1	GA Results	55
4.4.2	SNOPT Refinement	59
4.5	Cross-Case Comparison and Analysis	64
5	Conclusions	67
	Bibliography	68

List of Figures

2.1	Graphic representation of the lagrangain points in a generic system of 2 bodies. Credit to [3]	19
2.2	Graphic representation of the NRHO family Credit to [4]	20
2.3	Difference between single and multiple shooting. Credit to [5]	23
2.4	Graphic visualization for the difference between local/global minima and maxima. Credit to [7]	24
2.5	Hybrid optimization framework combining Genetic Algorithm (GA) for global exploration with Sequential Quadratic Programming (SQP) for local refinement. The GA identifies promising candidate trajectories, which are then precisely optimized by SQP using SNOPT.	26
2.6	Conceptual illustration of forward and backward propagation in Earth–Moon transfers. Forward propagation starts from LEO and integrates toward the NRHO, while backward propagation begins from the NRHO and traces trajectories back to Earth.	27
3.1	ECI reference frame	31
3.2	LVLH reference frame. Credit to [10]	32
3.3	Difference between the reference frames. Credit to [11]	32
3.4	Path relinking illustration. Credit to [13]	38
3.5	Weak stability boundary transfer. Credit to [2]	40
4.1	Flow chart of the the tool	42
4.2	Velocity around NRHO	46
4.3	Trajectory from different point of the NRHO	47
4.4	Legend of figure 4.3	47
4.5	Zoom of the arrival trajectories	48
4.6	Trajectory in ECI reference system with multiple shooting method adopted	49
4.7	Inclination trend w.r.t Earth during transfer	50
4.8	Components of velocity in RTN reference system for each maneuver	50
4.9	Keplerian orbital elements in the proximity of the Earth	51
4.10	Trajectory in ECI reference system after refinement	52
4.11	Inclination trend during the transfer refined	53

4.12	RTN components of velocity for the refinement case	53
4.13	Insertion into NRHO, view in IAU MOON reference system	54
4.14	Keplerian orbital elements in the proximity of the Earth	54
4.15	WBS transfer in ECI reference frame	56
4.16	Components of velocity in the wbs transfer	57
4.17	Inclination plane trend during the transfer	58
4.18	Parameter in the proximity of the Earth	59
4.19	Transfer in ECI reference frame refined with SNOPT	60
4.20	Zoom on the insertion point in the NRHO in ECI reference frame . .	61
4.21	Insertion maneuver in IAU MOON reference frame	62
4.22	Components of velocity in the wbs transfer	62
4.23	Inclination plane change trend during the transfer	63
4.24	Zoom departure from Earth	64

Acronyms

BLT ballistic lunar transfer. 12, 15, 16

BVP boundary value problems. 22

CR3BP Circular Restricted Three-Body Problem. 4, 16

ECI Earth Centered Inertial. 30

ESOC European Space Operations Centre. 4, 27

GA Genetic Algorithm. 4, 24

ISS International Space Station. 15

LEO Low Earth Orbit. 4, 12, 13, 15

NLP non linear programming. 22, 38

NRHO Near Rectilinear Halo Orbit. 4, 12, 13, 15

PMP Pontryagin Maximum Principle. 22

SPICE Spacecraft, Planet, Instrument, C-matrix, Events. 4, 29

SQP sequential quadratic programming. 4, 13, 24, 25, 34

WSB Weak Stability Boundary. 12, 15, 16, 39

Introduction

This thesis addresses the problem of designing and optimising low-energy and high-energy transfers between Earth and cislunar space, focusing specifically on connections between LEO and NRHO. Through analysing and comparing various transfer architectures and developing advanced optimisation techniques to minimise the necessary propellant while satisfying mission constraints.

The objectives of this thesis are focused on enabling efficient, reliable, and sustainable access to cis-lunar space, with a specific focus on transfers between LEO and NRHO around the Earth-Moon system. While the work is rooted in the theoretical framework of astrodynamics and optimal control, it is also strongly motivated by the practical requirements of upcoming missions associated with the Lunar Gateway. The objectives are therefore structured along three main axes: the primary optimization goal, related to minimizing the propulsive cost; the mission-level constraints, which define the operational context of the problem; and the additional spacecraft-related constraints, which ensure the feasibility of the proposed transfer architectures. Together, these elements form the foundation of the research and guide the methodological choices developed in the following chapters.

The principal goal of this thesis is the minimization of the total Δv required for the transfer trajectory. Since propellant mass is a dominant driver in spacecraft design, any reduction in Δv has a direct impact on the mission's performance, enabling either a reduction in launch mass, an increase in delivered payload, or an extension of the operational lifetime of the spacecraft. In this context, where repeatability of the mission is a crucial requirements, reducing the total amount of Δv and for directly consequence the total propellant became the first objective. In order to achieve this goal, is necessary to explore the solution space of the Sun-Earth-Moon system trough optimization techniques. Navigate the low-energy transfers solution, such as ballistic lunar transfer (BLT) and Weak Stability Boundary (WSB) trajectories, became essential in order to achieve a small consumption of propellant. This sensitivity highlights the need to implement robust numerical methods that can explore large solution spaces and converge towards feasible, high-quality trajectories.

The reference mission scenario under investigation is defined by transfers between two well-characterized orbital environments: a NRHO around the Earth-Moon system and a LEO. NRHO have been selected by international space agencies as the baseline operational orbit for the Lunar Gateway due to its unique dynamical prop-

erties. These include long-term stability, which implies minimal use of ΔV for station-keeping manoeuvres, a continuous line of sight with the Earth's ground station and regular access to the lunar surface. At the opposite end of the transfer, the target LEO is constrained by specific altitude and inclination values, chosen to reflect realistic launch and recovery conditions consistent with current launcher performance and reentry safety requirements. These constraints impose strict boundary conditions on the trajectory optimization problem. The transfer trajectory must not only achieve precise insertion into and departure from NRHO but also satisfy the geometrical and dynamical requirements of the LEO target. This dual-orbit configuration therefore defines a highly constrained design space in which the optimization methodology must operate, ensuring that all proposed solutions remain relevant for practical mission applications.

In addition to the objective of Δv minimization and the mission-level constraints associated with NRHO and LEO, the trajectory design problem must incorporate operational constraints related to spacecraft capability. These constraints include: limitations on the total propellant budget available due to a maximum tank dimension, limitation due to the total thrust generated by the engine, and restrictions on the time between two consecutive correction maneuvers, as well as, its total number. Such constraints are essential to ensure that the resulting trajectories are not only mathematically optimal but can also be executed by spacecraft systems available today or in the near future. Moreover, an optimization process that minimizes Δv but requires long duration transfers cannot be suitable with crewed missions or time-critical cargo deliveries. This additional constraints need to be implemented directly into the optimization problem, rather than treating them as post-processing filters. This integrated approach ensures that the trajectories identified are robust, practical, and aligned with mission requirements.

This thesis is organised into five main chapters, providing a coherent progression from the initial analysis to the final result and guiding the reader through the work done.

Chapter 1 introduces the overall mission context, emphasizing the growing importance of cis-lunar space and the role of the Lunar Gateway as a strategic infrastructure for future exploration.

Chapter 2 provides the theoretical background and a review of the state of the art. It covers the fundamental concepts of astrodynamics, including the restricted three-body problem and the characterization of NRHO orbits, as well as the mathematical foundations of trajectory optimization. It also includes specific algorithms employed in this work, with a focus on global methods such as Genetic Algorithms and local approaches such as SQP.

Chapter 3 describes the methodology and problem formulation. The high-fidelity dynamical model adopted for the simulations is introduced, along with the use of ephemerides for accurate propagation. The optimization problem is then formally defined, detailing the decision variables, objective function, and constraints. The hy-

brid optimization strategy, which combines global exploration with local refinement, is explained in depth. In addition, advanced strategies such as multi-start searches, launch date variation, and weak stability boundary trajectories are presented as extensions to broaden the exploration of the solution space.

Chapter 4 presents the implementation, results, and analysis. It presents the tool logic flow, and the explanation of how the optimisation problem was implemented in the two cases of global and local searches, then all the results are presented. Initial solutions identified with the Genetic Algorithm are compared with the refined solutions obtained through SQP, highlighting improvements in convergence and efficiency. The optimized trajectories are analyzed in detail, with graphical representations in different reference frames, as well as discussions on maneuver distribution, sensitivity, and the role of advanced exploration strategies.

Finally, Chapter 5 summarises the conclusions of this work and explains the strengths and weaknesses of the approach applied to obtain the results. In addition, it suggests some future developments.

1. Mission context

In this section will be presented the motivation behind this work and provide an overview of upcoming missions in the cislunar domain.

1.1 Access to cislunar space and the lunar Gateway

There is an incremental interest in the space between the Earth and the Moon, the so-called 'cis-lunar space'. This region of space has become a central objective of current international space programmes, carrying a significant shift in focus away from traditional LEO. The cislunar region is increasingly being seen as a vital environment for enabling sustainable space exploration. This region is important for this reason: it is an accessible space in which to test new technologies; it enables the establishment of hubs for operations on the lunar surface; it constitutes a natural gateway to distant destinations, including Mars and beyond.

From this perspective, the Lunar Gateway is a significant milestone in infrastructure development. Designed as a modular station to be serviced by crews and positioned in a near-rectilinear halo orbit (NRHO) (see e.g., [1]) around the Earth–Moon L2 lagrangian point, it is being developed as part of NASA's Artemis programme. The Gateway is designed to fulfil a variety of roles. It will serve as a platform for lunar landings, a laboratory for conducting scientific and technological experiments in different conditions from those experienced on the International Space Station (ISS). Furthermore, its NRHO location provides unique operational advantages by combining long-term orbital stability and a continuous line of sight to both Earth and the lunar surface.

Despite these advantages, establishing a permanent presence in cis-lunar space presents not simple mission design challenges. Unlike traditional LEO operations, transfers to the NRHO require careful consideration of multi-body dynamics, long-duration mission constraints, and limited propellant availability. Therefore, in terms of propellant consumption, efficient access to the Gateway must be guided by advanced trajectory design methodologies capable of exploiting low-energy pathways, such as BLT and WSB (see e.g., [2]) trajectories. These require a long mission duration and low fuel consumption, making them optimal for cargo missions but not crewed ones.

1.2 The low-energy transfer problem

Designing efficient trajectories between Earth and the Moon has long been a central challenge in astrodynamics. Traditional approaches rely on direct transfers, which, although are short in duration, typically require high propellant consumption. This makes them less attractive for missions constrained for the high request of fuel budgets. In contrast, low-energy transfers exploit the dynamical environment of the Earth-Moon-Sun system to reduce the overall cost in terms of Δv . These transfers exploit the intrinsic properties of multi-body dynamics, in which gravitational perturbations from the Moon and Sun can be used in order to achieve propellant saving.

The common dynamic model employed for the analysis of some trajectories is the circular restricted 3 body problem (CR3BP), which provides valuable insights into the underlying dynamics and the role of invariant manifolds in guiding spacecraft along energy-efficient paths. Extensions to higher-fidelity N-body models allow for the inclusion of additional perturbations, such as solar gravity, solar radiation pressure and the non-spherical gravitational fields (harmonic perturbations) of the Earth and Moon. Within this context, families of trajectories such as BLT and WSB trajectories have been analyzed. These paths are characterized by longer flight times compared to direct transfers, but they offer significant reductions in propellant consumption, making them particularly appealing for cargo missions, resupply operations, and exploratory missions.

The main challenge of low-energy transfers lies in the complexity of the optimisation problem. In fact, low-energy transfers are characterised by their sensitivity to initial conditions and require more challenging trajectory design.

This thesis addresses the issue of high- and low-energy transfer by taking a hybrid optimisation approach, combining global and local solvers to achieve optimal results. This work is necessary to identify feasible solutions or trajectories that minimise fuel requirements while respecting the imposed constraints. This is important for future missions to the Lunar Gateway.

2. Theoretical background and state of the art

2.1 Fundamentals of astrodynamics

Astrodynamics is essentially the field of space science that looks at how objects, both natural and artificial, move when they are influenced by gravitational and, in many cases, non-gravitational forces coming from celestial bodies. In simple terms, it is about understanding and predicting how satellites, spacecraft, or even smaller fragments of matter travel through space. This discipline provides not only the theoretical background but also the computational tools needed to design and analyze trajectories. Thanks to these methods, it becomes possible to model satellite orbits with good accuracy and to plan more complex missions, such as interplanetary transfers. In practice, astrodynamics connects abstract mathematics to very concrete mission needs.

The following sections will focus on the dynamical models that govern the motion of objects in space. The discussion begins with the most basic models and approaches, which still capture the essential features of orbital motion and provide an initial understanding of how bodies move. Although these dynamical models are simplified, they provide an important foundation for understanding the physics of the problem and possible solutions. After this initial analysis, the model will be complicated by adding additional perturbations, such as solar radiation pressure or the gravitational harmonics of a celestial body. This makes the results of the study more realistic.

2.1.1 The two-body problem

The two-body problem offers the simplest yet most fundamental description of orbital motion. Assuming that only two bodies interact through mutual gravitation, the trajectory of a spacecraft is governed by Kepler's laws and takes the form of a conic section, such as an ellipse, parabola, or hyperbola. This motion satisfies Newton's law of gravitation.

$$\ddot{\mathbf{r}} = -\mu \frac{\mathbf{r}}{r^3} \quad (2.1)$$

where $\mathbf{r}$ is the position vector of the spacecraft with respect to the central body, $r = \|\mathbf{r}\|$, and $\mu = G(M+m)$ is the standard gravitational parameter. The specific mechanical energy of the orbit is conserved and can be expressed as

$$\epsilon = \frac{v^2}{2} - \frac{\mu}{r} \quad (2.2)$$

with v being the velocity magnitude. The orbital shape is characterized by the semi-major axis a , related to the energy through

$$a = -\frac{\mu}{2\epsilon} \quad (2.3)$$

The position and motion of an object in orbit can be captured using six classical, or Keplerian, elements. These parameters make it possible to represent the orbit in a way that is both compact and intuitive, turning a complex three-dimensional motion into a clear set of numbers. Even though this formulation is relatively simple and leaves out many of the real-world perturbations that affect spacecraft, it continues to be extremely useful. In fact, it is still the standard choice for preliminary mission studies, quick trajectory checks, and for defining essential orbital properties such as the orbital period, the shape of the path through eccentricity, and the tilt of the orbit through inclination. Its lasting relevance comes from the balance between simplicity and the immediate physical insight it provides.

2.1.2 The restricted three-body problem (CR3BP)

Although the two-body framework offers a clear and elegant way to describe orbital motion, it quickly shows its limits when dealing with environments where more than one massive body plays a significant role. A prime example is the Earth–Moon system, where the gravitational force of both bodies must be considered. To address this, the restricted three-body problem (CR3BP) extends the classic model by including two primaries—such as Earth and Moon—while treating the spacecraft as a third body with negligible mass.

What makes this model particularly interesting is that it reveals dynamical behaviours that cannot be captured by the simpler two-body description. Among the most important are the five Lagrange points, which correspond to equilibrium solutions in the rotating reference frame. Around these points, one can find families of periodic and quasi-periodic orbits, such as halo orbits and Lissajous trajectories, which have received considerable attention in mission design. Another key feature of the CR3BP is the existence of invariant manifolds. Thanks to these pathways, the spacecraft was able to navigate space without using any propellant. Analysing this model makes it possible to identify promising trajectories that could save a significant amount of propellant, which is a particularly valuable advantage in the context of long-duration or logistics missions in the Earth–Moon system.

2.1.3 Lagrange points and periodic orbits

The five Lagrange points of the Earth-Moon system play a central role in trajectory design for cis-lunar missions. They correspond to equilibrium configurations where the gravitational forces of Earth and Moon balance the centrifugal force in the rotating reference frame. As shown in 2.1, the collinear points (L_1 , L_2 , L_3) are located along the Earth–Moon axis. L_1 and L_2 are at the same distance from the secondary body. L_3 is located in the opposite direction to the secondary body, at the same distance from the primary as the secondary. The points L_4 and L_5 form an equilateral triangle with the primary. . Around these points, various families of periodic and quasi-periodic orbits exist, such as halo orbits and Lissajous trajectories. These orbits are of great interest because they provide staging locations for spacecraft and allow efficient connections to other regions of the cis-lunar space by exploiting invariant manifolds.




Figure 2.1: Graphic representation of the lagrangain points in a generic system of 2 bodies. Credit to [3]

2.1.4 Near-Rectilinear Halo Orbits (NRHO)

Among the different periodic orbits that can be found around the Earth–Moon L_1 and L_2 points, the so-called Near Rectilinear Halo Orbits (NRHOs) have gained particular attention in recent years. These orbits are quite peculiar. They are long and stretched, and almost rectilinear in shape. If a spacecraft flight takes place on one of these orbits, then most of its trajectory will be rectilinear. In addition, a peculiarity is that the spacecraft flies very close to the lunar poles at one point and then sends it far away from the surface at the other. This geometry is unusual but very useful. In fact, NRHOs bring together several advantages that are not easy to find all in one orbit. They keep a stable line of sight to Earth, which makes communication reliable; they do not demand excessive amounts of ΔV for station-keeping, which is always a critical factor; and they create repeated opportunities to pass close to the Moon, which naturally helps with both crew and cargo transfers. This makes the study of NRHOs not just interesting in theory, but absolutely central for the kind of trajectory optimization problems considered in this thesis.




Figure 2.2: Graphic representation of the NRHO family Credit to [4]

Figure 2.2 shows the family of NRHOs, which is subdivided into four categories: L_1 , L_2 northern and L_1 , L_2 southern. The specific orbit considered in this study is a southern NRHO around the Earth–Moon L_2 point, set in a 9:2 resonance with the Moon’s motion. In simple terms, this resonance indicates that for every two revolutions of the Moon around the Earth, the spacecraft manages to perform nine complete loops of its own. The practical outcome is an orbital period of about six and a half days, which is long enough to allow for planning and operations, yet short

enough to maintain regular opportunities for transfer.

2.1.5 Perturbations and N-Body models

Even though the CR3BP does a pretty good job at capturing the main features of cis-lunar dynamics, it is still, a simplified or idealized model. In reality, if one wants to plan missions with higher fidelity in the Sun–Earth–Moon system, more effects need to be considered. For example, there are gravitational influences from the Sun and other planets, higher-order terms in the Earth’s and Moon’s gravity fields, and even non-gravitational forces like solar radiation pressure. Each of these effects might be small on its own, but over long mission durations, they can add up and start to noticeably change the spacecraft’s trajectory. So, to account for these influences, it is common to use a full n-body model combined with numerical integration. Of course, this approach gives a more realistic picture of what the spacecraft will actually do, but it also comes with a trade-off: the computations become more complicated, and simulations take longer to run. In practice, this is the reason why careful balance between model accuracy and computational cost is always necessary in trajectory design.

2.1.6 Role in mission design

Models such as the two-body problem or the CR3BP are useful for gaining an initial understanding of the problem and identifying the main mechanisms and promising families of trajectories. They are simple, easy to use, and don’t require heavy computations, which makes them perfect for initial analyses or for just getting a feel for the problem. But of course, these simplified models have a limit. When it comes to actually implementing a mission and making sure everything will work in practice, it is necessary to get more accurate results. That’s where numerical simulations come in, especially n-body models, which can capture the combined effects of multiple gravitational bodies and other forces. Using these high-fidelity models gives a much better picture of what the spacecraft will really do, even though they are more computationally demanding. In short, the combination of simple models for insight and complex models for precision is really what makes trajectory design possible in the real world.

In this study, the use of a high-fidelity model ensures that the results obtained simulate the real physics of the problem. Additionally, this study enables us to understand how the spacecraft moves during its trajectory. Although this model requires a lot of computing time, it gives this study the right level of importance and provides confidence that the trajectories are not only theoretical, but can also be applied in practice.

2.2 Trajectory optimization methods

Designing optimal trajectories is one of the trickiest challenges in astrodynamics, especially when dealing with multi-body environments like the Earth–Moon–Sun system. The dynamics in these regions are extremely sensitive, so even small changes in initial conditions can lead to very different outcomes. On top of that, there are mission constraints to consider, which makes the problem even more delicate. Because of this, it's necessary to rely on solid mathematical frameworks that can both formulate and solve the optimization problems effectively. Trajectory optimization methods fall into two main categories: direct approaches and indirect approaches. Among the direct methods, shooting techniques — particularly single shooting and multiple shooting — are widely used because they help maintain convergence and allow better control over the trajectory. In practice, choosing the right method often depends on the specific problem at hand and the level of precision required, as well as on how sensitive the trajectory is to small perturbations.

2.2.1 Direct vs. Indirect Approaches

Trajectory optimization can be approached in a couple of different ways, either by directly transcribing the problem into a numerical format or by using indirect methods rooted in the calculus of variations. Indirect approaches find the optimality through the Pontryagin Maximum Principle (PMP) or by using Hamiltonian mechanics, which ends up producing a set of boundary value problems (BVP). In practice, this method tends to be very sensitive to the initial guesses and can become quite challenging when dealing with complex problems, constraints, or additional perturbations.

On the other hand, direct approaches use a more pragmatic approach. The trajectory is discretized into a finite set of variables, for example, impulses, and the problem is transformed into a non-linear programming (NLP) formulation. This allows the use of modern numerical solvers that can resolve large-scale problems with multiple constraints. Direct methods introduce some approximations, but they are generally more robust and flexible, which makes them much better suited for real mission design. This is especially true in the cis-lunar environment, where the dynamics are influenced by multiple bodies and a range of perturbations, and where practical, implementable solutions are more valuable than mathematically exact ones.

2.2.2 Shooting techniques: single vs. multiple shooting

Within the problem of direct optimization, shooting techniques are a widely used approach to maintain continuity of the trajectory while optimizing control inputs. In the single shooting method, the entire trajectory is propagated starting from

a single initial condition, and the optimization variables are applied only at that starting point.

The method is simple to implement and provides a good starting point for an initial feasibility study. However, it has the weakness of concentrating the adjustment at the beginning, which makes it very sensitive to initial guesses. Small errors can grow over time, adding numerical instability to the solution. This is particularly important with long-time propagation and high nonlinear dynamical models, such as those applied in this study.




Figure 2.3: Difference between single and multiple shooting. Credit to [5]

To overcome these issues, as it is possible to see in figure 2.3 the multiple shooting method (see e.g. [6]) divides the trajectory into several segments, introducing intermediate nodes at which additional decision variables are applied. The trajectory is then reconstructed by enforcing continuity constraints between the segments. This significantly improves the controllability of the optimization problem, reduces sensitivity to local instabilities, and allows for a finer adjustment of maneuver distribution along the path. On the other hand, this method leads to an increase in the dimensionality of the problem, which requires more computational resources and careful calibration of the number of nodes.

When dealing with low-energy transfers in the cis-lunar environment, the multiple shooting technique often turns out to be far more reliable than the single-node approach. This structure makes it possible to respect orbital requirements — like maintaining altitude, hitting the right inclination, or matching insertion conditions — without the solution drifting away. Another strength of this method is its flexibility: because the trajectory is partitioned, it becomes easier to include further operational limits, whether they relate to available propellant, maximum thrust levels, or even launch and maneuver timing windows. In practice, this allows for trajectories

that are not only efficient in terms of ΔV , but also realistic and consistent with what an actual spacecraft could carry out.

Multiple shooting, on the other hand, adds a growing number of decision variables for each node applied. This makes the trajectory more robust, but increases the total amount of computational time. In addition, by combining Multiple Shooting with GA and SQP, the method achieves greater robustness, enhanced convergence, and improved ΔV efficiency, making it particularly suitable for the transfers under analysis.

2.3 Optimization algorithms

Trajectory optimization in astrodynamics usually comes down to finding good numerical methods that can handle large and complicated search spaces, while at the same time respecting all the mission requirements. The process is rarely straightforward, because an algorithm has to explore enough of the solution space to avoid missing good options, but it also needs to converge with precision once a promising path is found. In fact, looking at the figure 2.4, it is possible to appreciate the difference between a global minimum and a local minimum. In general, a global optimizer like the GA generally leads to a global minimum because it cannot see local minima due to its low accuracy. Conversely, a local optimiser can see a path through to the local minimum.

For this reason, most approaches rely on the hybrid approach combining global exploration that helps map out the broad regions where viable solutions might exist, and local refinement then takes over to fine-tune the trajectory.




Figure 2.4: Graphic visualization for the difference between local/global minima and maxima. Credit to [7]

2.3.1 Global optimization: genetic algorithms (GA)

Genetic Algorithms are part of the broader family of stochastic, population-based optimization methods, loosely inspired by the way evolution works in nature. The basic idea is that instead of working with a single candidate solution, a whole population of possible trajectories is generated and then gradually improved over several generations. At each step, genetic operators such as selection, crossover, and mutation are applied, then natural selection keep the better solutions and mixing their features to create new ones.

The quality, or “fitness,” is measured by measuring the sum of the objective function and the constraints violation. In trajectory optimization problems, the objective function is often related to the total ΔV required. In this way, the algorithm explores a wide range of potential solutions without being tied to a single starting point, making it particularly useful in complex environments like cis-lunar transfers, where the search space is full of local minima and small variations can lead to very different outcomes.

2.3.2 Local optimization: sequential quadratic programming (SQP) and the SNOPT solver

While global optimization methods are very good at scanning all the solution space and finding out regions where something interesting might lie, they usually need to be supported by local optimization techniques in order to refine those rough candidates into precise, usable trajectories. One of the most used approaches in this category is SQP, a standard tool for constrained nonlinear optimization problems.

Among the various solvers that implement this approach, SNOPT has proven particularly effective. It is designed to handle large-scale optimization problems where many variables and constraints are involved, which makes it a natural choice for trajectory design. In practice, this means it can deal with path constraints, boundary conditions, ΔV limits, insertion requirements, and other operational restrictions that are essential in real mission planning.

2.4 Forward vs. backward propagation approaches

In trajectory design, the choice of propagation direction represents a critical methodological aspect. The most intuitive approach is forward propagation, in which the spacecraft dynamics are integrated forward in time starting from well-defined initial conditions. This method is commonly adopted in mission analysis, as it provides a direct correspondence with the physical timeline of the mission and allows for the straightforward evaluation of state evolution under given control inputs. However, forward propagation can become inefficient or even impractical when the terminal conditions are of primary importance, as in the case of transfers to periodic orbits or




Figure 2.5: Hybrid optimization framework combining Genetic Algorithm (GA) for global exploration with Sequential Quadratic Programming (SQP) for local refinement. The GA identifies promising candidate trajectories, which are then precisely optimized by SQP using SNOPT.

libration point trajectories. In such cases, small deviations in the initial guess can lead to large mismatches at the final state due to the sensitivity of the dynamical system.

Backward propagation offers a complementary perspective. Here, the integration is performed backward in time, starting from the target conditions and propagating toward the departure region. This approach is useful when the target orbit is difficult to reach; instead, the starting orbit is easy to reach; for instance, when the orbital elements of a target orbit (such as an NRHO) are precisely known from ephemerides. By propagating backward, one can generate families of candidate trajectories that are naturally consistent with the final constraints, significantly reducing the search space of admissible solutions. Furthermore, backward propagation facilitates the identification of insertion points and transfer opportunities that would be more challenging to detect with forward integration alone.

In the present work, backward propagation was selected as the primary approach for trajectory exploration and optimization. Starting from reference states of the NRHO provided by ESOC ephemerides, trajectories were propagated backward in time to identify feasible connections with low Earth orbits. This methodology not only improved the robustness of the optimization process but also aligned with the objective of minimizing ΔV while satisfying mission constraints.




Figure 2.6: Conceptual illustration of forward and backward propagation in Earth–Moon transfers. Forward propagation starts from LEO and integrates toward the NRHO, while backward propagation begins from the NRHO and traces trajectories back to Earth.

3. Methodology and problem formulation

This chapter gives a panoramic view of the design and optimization of transfer trajectories from a Near-Rectilinear Halo Orbit (NRHO) around the Moon to a Low Earth Orbit (LEO). The dynamic model used, the optimisation technique employed, and the strategy for identifying optimal solutions will be presented.

The approach used aims to obtain solutions, i.e. transfers, at low cost in terms of Δv . This is achieved by using advanced optimisation techniques and high-fidelity dynamic models, which faithfully simulate the physical nature of the problem.

3.1 High-fidelity dynamical model

To achieve realistic results that closely reflect the physics of the problem, a dynamic n-body model was implemented. In particular, the Earth, Moon, and Sun were taken into consideration; in fact it is especially important for transfers that exploit weak stability boundaries.

Unlike the two-body model, the n-body model allows the gravitational action of any massive body in the Solar System to be considered. Clearly, the smaller the mass of a body and the farther it is from the spacecraft under analysis, the less influence it will have on the spacecraft's dynamics.

The equations of motion for the n -body problem are expressed as:

$$\ddot{\mathbf{r}}_{sc} = -\frac{\mu_{\text{Earth}}\mathbf{r}_{sc}}{|\mathbf{r}_{sc}|^3} + \sum_{j=2}^n \mu_j \left(\frac{\mathbf{r}_j - \mathbf{r}_{sc}}{|\mathbf{r}_j - \mathbf{r}_{sc}|^3} - \frac{\mathbf{r}_j}{|\mathbf{r}_j|^3} \right) \quad (3.1)$$

where

- $\ddot{\mathbf{r}}_{sc}$: total acceleration vector of the spacecraft in the Earth-Centered Inertial (ECI) frame [km/s²].
- μ_{Earth} : Earth's gravitational parameter.
- $\mathbf{r}_{sc}$: position vector of the spacecraft relative to Earth's center in ECI coordinates [km].
- $\sum_{j=2}^n$: summation over all n third-body perturbing celestial bodies (Sun, Moon, planets).

- μ_j : gravitational parameter of the j -th perturbing body [km^3/s^2].
- $\mathbf{r}_j$: position vector of the j -th body relative to Earth's center in ECI coordinates [km].
- $\mathbf{r}_j - \mathbf{r}_{sc}$: vector from spacecraft to the j -th body [km].
- $|\mathbf{r}_j - \mathbf{r}_{sc}|^3$: cube of the distance between spacecraft and j -th body [km^3].
- $\frac{\mathbf{r}_j - \mathbf{r}_{sc}}{|\mathbf{r}_j - \mathbf{r}_{sc}|^3}$: direct term - gravitational acceleration that the j -th body exerts directly on the spacecraft
- $\frac{\mathbf{r}_j}{|\mathbf{r}_j|^3}$: indirect term - correction for the non-inertial acceleration of Earth due to the j -th body, necessary because ECI frame is not truly inertial

In this thesis, the set of perturbing bodies includes the Earth, the Moon ($j=2$), and the Sun ($j=3$). Their positions and velocities are evaluated from the SPICE ephemerides, ensuring a high degree of accuracy in the dynamical representation.

3.1.1 N-Body propagation and numerical integration

Building on the adopted n-body formulation, the numerical propagation proceeds segment-by-segment with an adaptive high-order Runge–Kutta integrator, so that local error control preserves fidelity over long cislunar arcs while keeping runtime manageable in the subsequent multiple-shooting solve. At each segment boundary, impulsive maneuvers are applied as instantaneous velocity updates and the integration is restarted. Event handling is used to trigger impulses and check path conditions such as altitude and safety margins, ensuring segment integration remains compatible with mission constraints rather than correcting violations only in post-processing. Time-scale and frame choices are kept consistent with those used to generate states for the primaries, preventing phase drift between integrated dynamics and ephemeris-driven positions across multi-weeks transfers and preparing a smooth hand-off to the ephemerides description in the next subsection. This propagation setup is thus intentionally designed to reduce long-arc sensitivity before optimization while providing the structure (nodes, defects, events) that multiple shooting exploits to converge reliably under operational constraints.

3.1.2 Use of SPICE ephemerides (JPL/ESOC)

This work relies on SPICE (see e.g., [8]) to supply time-tagged states and ancillary data with consistent frames and time systems, ensuring that n-body accelerations are computed against authoritative Earth, Moon, and Sun ephemerides at the exact propagation epochs. A standard kernel set is assumed:

- SPK for planetary/lunar/solar states


- PCK for body orientation and radii as needed for altitude and safety checks
- LSK to manage conversions between UTC and ephemeris time (ET/TDB) so integration timestamps and kernel queries remain synchronized

Frame handling is explicit and conservative: all force evaluations and state propagation occur in the inertial J2000/EME2000 (see e.g., [9]) frame to keep the ODE right-hand side free from per-step rotations, avoiding numerical noise and drift introduced by mixed-frame dynamics. Transformations to analysis frames (e.g., Moon-centered body-fixed for altitude checks or RTN for maneuver bookkeeping) are applied only after propagation, so the integrator always “sees” a consistent inertial geometry across segments. Each multiple-shooting node stores its ephemeris time (ET/TDB) alongside the state; integration advances ET monotonically and node ET stamps are reused verbatim during optimizer re-evaluations to guarantee deterministic impulse application and identical ephemeris queries when recomputing defects and gradients. At every dynamics call, Earth, Moon, and Sun states are interpolated at the current ET directly in J2000, eliminating intermediate transforms and ensuring the n-body accelerations reflect the true geometry at that epoch; immediately after an impulsive Δv , the velocity is updated and integration restarts from the same ET to preserve a one-to-one mapping between nodes and ephemeris samples. The NRHO target is represented as an ephemeris-parameterized path indexed by ET in J2000, so the insertion condition reduces to selecting a scalar (or small tuple) that locates an insertion epoch and state on the reference orbit; sampling by ET enables backward propagation from precise insertion points and shrinks the search space while keeping terminal conditions anchored to a single, consistent SPICE-based description of the NRHO.

3.1.3 Reference frames and usage

This section summarizes the reference frames adopted in the study, the reason for their selection, and the specific roles they play within propagation, optimization, and analysis. Frame definitions are kept consistent across all stages to ensure that state vectors, constraints, and visualizations are mutually coherent.

Inertial, body-fixed, and local frames. As it shown in figure 3.1 the Earth Centered Inertial (ECI) is a frame aligned with the mean equator and equinox at J2000 (e.g., J2000/ICRF), used as the default propagation and optimization frame. The Z -axis is normal to the mean equatorial plane at epoch J2000; the X -axis points toward the mean vernal equinox; the Y -axis completes a right-handed triad. This frame avoids artificial accelerations due to Earth rotation and is therefore used for dynamics and constraint enforcement on state continuity and Δv minimization.

**Figure 3.1:** ECI reference frame

Moon-centric frames: A Moon-centered frame is employed when sampling or analyzing NRHO states and Moon-relative geometry (*e.g.*, for distance filtering and insertion phase characterization). Depending on the task, a Moon inertial frame (*e.g.*, Moon-J2000) or a body-fixed frame (*e.g.*, IAU_MOON) may be referenced; in this work, inertial sampling is preferred for trajectory sensitivity and phase selection, while body-fixed quantities are used only when surface-referenced metrics are required.

Local orbital frames (LVLH/VNB): For interpreting maneuver geometry or reporting local kinematics, local frames attached to the spacecraft are optionally used. LVLH aligns Z toward nadir and X along the velocity projection (local horizontal), whereas VNB aligns axes with velocity (V), orbit normal (N), and binormal (B). These frames are visualization/diagnostic tools and are not employed for propagation.




Figure 3.2: LVLH reference frame. Credit to [10]

In this work, employing both an Earth-Centered Inertial (ECI) and a Moon-centered reference frame is essential to characterize the geometry of the NRHO, as each frame yields a distinct and complementary visualization of the orbit. In the following images, figure 3.3, it is possible to appreciate this difference.




Figure 3.3: Difference between the reference frames. Credit to [11]

Frame usage in the pipeline. Propagation and optimization: all trajectory integrations, multiple-shooting continuity conditions, and the nonlinear program (objective and constraints) are expressed in the ECI frame.

Constraint evaluation: Orbital inclination, altitude, and time-of-flight are evaluated from inertial states to ensure compatibility with the optimization transcription. When Earth-fixed constraints are required, inertial states are mapped to the Earth-fixed frame for the specific check, and results are reported back in the inertial analysis context.

Maneuver representation: Impulsive maneuvers are stored and applied in the inertial frame to guarantee consistency with the integrated dynamics and with the multiple-shooting defect constraints. When needed for interpretation, impulses may be projected into LVLH/VNB at the node epoch without affecting the optimization.

Transformations, epochs, and units. Frame transformations: all transformations between inertial, body-fixed, and local frames follow the frame chain defined by the adopted kernels/frame definitions. Transformations are applied at the exact epoch of interest to avoid temporal interpolation inconsistencies.

Time scales: Ephemeris access and transformations are performed at consistent epochs (ET/TDB or the chosen internal time scale). Timestamps reported in figures and tables indicate the time scale used. Any conversions to UTC for readability are applied only at the presentation layer.

Units and constants: States are expressed in kilometers and kilometers per second; gravitational parameters, radii, and derived quantities are used consistently in these units. Geodetic quantities (when reported) are explicitly labeled and distinguished from geometric radii used in inertial-space altitude proxies.

3.2 Formulation of the optimization problem

The optimization problem is posed as a constrained nonlinear program over a multiple-shooting transcription that minimizes total impulsive cost while satisfying dynamical, path, and operational constraints between a specified NRHO insertion state and a target LEO configuration. Decision variables comprise the sequence of impulsive maneuvers $\Delta \mathbf{v}_i$ at segment boundaries, segment durations τ_i , the departure epoch, and a compact parameter selecting the insertion point along the ephemeris-based NRHO reference, enabling the solver to adjust maneuver placement and timing while anchoring terminal conditions to a consistent target path. The objective function is the sum of maneuver magnitudes.

$$J = \sum_{i=1}^N \|\Delta \mathbf{v}_i\| \quad (3.2)$$

directly correlating with propellant through Tsiolkovsky and standard for impulsive transfer design in cislunar mission analysis. Dynamics are enforced segment-wise

by integrating the n-body model with ephemeris-driven primaries; the continuity “defect” between the post-impulse state of segment i and the initial state of segment $i+1$ is constrained to zero within tight tolerances to ensure a physically continuous trajectory across all nodes. Terminal constraints impose the LEO geometry (altitude range and inclination) at Earth return and an insertion tolerance about the selected NRHO state (position/velocity bounds), while path constraints bound Earth/Moon altitudes, enforce flyby safety margins, limit maximum impulse magnitudes, and impose a minimum time separation between maneuvers to avoid clustering. The formulation is solved in a hybrid workflow where a global GA explores departure dates, insertion parameters, maneuver counts, and coarse segment durations to provide feasible seeds, and a local SQP (SNOPT) refines impulses.

3.2.1 Decision vector (Delta-V and flight times τ)

The decision vector collects all control and timing quantities required by the multiple shooting transcription: impulsive maneuvers Δv_i applied at the boundaries of N segments, the segment flight times τ_i , the mission departure epoch t_0 , and a compact insertion parameter s that selects the target state along the ephemeris-based NRHO reference path. Each Δv_i is a three-component vector in the inertial propagation frame, bounded in magnitude to reflect propulsion capability and constrained by a minimum inter-impulse time enforced through the τ_i variables; the τ_i themselves are lower-bounded to avoid maneuver clustering due to a technical engine limit and upper-bounded to limit total time of flight in line with mission class (cargo vs crew). The scalar s parametrizes the NRHO insertion state by ephemeris time (or an equivalent phase/arc-length variable), reducing terminal decision dimensionality while preserving precise targeting to a consistent reference orbit; together with t_0 , it controls phasing between Earth departure and lunar insertion without over-parameterizing the final boundary conditions. This structure enables the optimizer to redistribute maneuver magnitudes and timings, and satisfy all path and terminal constraints, while maintaining a compact and numerically well-conditioned search space for the local refinement.

3.2.2 Objective function: minimization of total Delta-V

The objective is to minimize the total impulsive cost, defined in the same form of equation 3.2, but its use differs across the two optimization stages: in the global search, the genetic algorithm treats J as a primary fitness component combined with constraint penalties (e.g., terminal/path violations and other constraint that will be presented in detail in the next paragraphs) to rank and evolve diverse candidates rather than to enforce tight feasibility; conversely, in the local refinement, the SQP solver minimizes the same J under explicit equality/inequality constraints with scaled tolerances, and violations toward zero while exploiting derivatives for precise

convergence. In practice, the GA’s penalty objective prioritizes broad exploration and robustness to poor initial guesses, accepting loosely feasible or mildly infeasible individuals if they improve Δv trends, whereas SNOPT operates on a strictly constrained problem where J is reduced only along directions that preserve feasibility to first order, ensuring accurate impulse sizing and timing once a credible seed is available.

3.2.3 Boundary and path constraints (altitude, inclination, safety)

Terminal and path constraints are tailored to the transfer class to guide the global search toward feasible architectures before local refinement. For direct transfers, terminal constraints enforce insertion into the selected NRHO state within tight position/velocity tolerances, selected from the SPICE’s ephemeris reference, and Earth return into a target LEO characterized by an altitude, inclination, and a window tied to launch site/access constraints, and bounded eccentricity to ensure operational circularization. Path constraints include safe altitudes with Earth to protect against atmospheric or terrain interception, with Moon to avoid fly-by solutions.

For WSB/low-energy transfers, terminal constraints at NRHO remain, but the path constraints expand to include apogee distance ranges beyond Earth distance (e.g., 1.2–1.5 million km), according to [12] and Sun–Earth geometry windows that realize the required energy shaping; additional constraints include Earth departure C3 and enforce low lunar arrival energy compatible with ballistic capture.

3.3 Hybrid optimization strategy (GA + SQP)

This work adopts a two-stage hybrid strategy that leverages global exploration to discover feasible transfer architectures and local refinement to achieve high-accuracy, constraint-satisfying solutions. In the first stage, a genetic algorithm (GA) explores departure epochs, NRHO insertion parameters, maneuver counts, and segment durations under graded constraint penalties, promoting diversity while gently steering the search toward low- Δv . In the second stage, a sequential quadratic programming (SQP) solver refines each candidate using the multiple-shooting transcription, driving continuity defects and terminal/path violations within tight tolerances while reducing total impulsive cost. By separating exploration and refinement, the framework attains low- Δv via the local solver and robust feasibility through globally guided constraint handling.

3.3.1 Phase 1: global search with genetic algorithm and multiple shooting

Global search (GA). The GA operates on a compact decision vector comprising:

- three components of velocity v_{iN} , where index i is related to the x,y,z components and N the number of nodes of the trajectory
- τ_j $N-1$ variables

. Each segment duration is parametrized by a dimensionless variable τ , scaled by reference time T_{ref} so that the physical duration is τT_{ref} ; the limiting case $\tau = 1$, corresponds to allocating the entire reference duration to a single segment, which recovers the single-impulse limit of the transfer model. Fitness combines the total Δv with penalty terms that quantify infeasibility, including terminal mismatches (NRHO/LEO), path-safety violations (altitude, approach geometry), time-of-flight overruns, and limit fly altitude for the Earth and Moon. In the GA stage, a weighted fitness is used: each objective or penalty term carries a coefficient that calibrates its influence on selection, ensuring that feasibility pressures and performance goals (low Δv) are appropriately balanced during evolution.

$$F = w_v v + w_{altitude} P_{altitude} + w_{inclination} P_{inclination} \dots \quad (3.3)$$

As shown in Equation 3.3, w_i controls the contribution of each objective/penalty term and P_i encodes constraint violations; together they shape the search landscape so the GA converges toward solutions that satisfy the imposed constraints.

Penalties are weighted in order to preserve promising but imperfect solutions, while “hard” penalties are reserved for safety violations. Typical settings are: population size 80, generation 400. The GA outputs one or more near-feasible candidates with competitive Δv and acceptable constraint residuals for local refinement.

3.3.2 Phase 2: local refinement

Each GA seed is transcribed into a multiple shooting NLP with decision variables Δv_i , τ_i , and the constraints that enter the GA as weighted penalty terms in the fitness are enforced explicitly as inequality constraints during the local refinement stage. The objective (eq. 3.2) subject to constraints, with scaling applied to states, and penalties to improve conditioning. The solver uses analytic or semi-analytic sensitivities where available; otherwise, finite-difference or automatic-differentiation approximations are applied with step sizes tuned to ephemeris/interpolator accuracy. Tolerances are set to 10^{-6} for constraints.

3.4 Advanced strategies for solution space exploration

After generating candidate seeds with the GA, each solution is refined by the local solver; however, because the search landscape is highly multimodal with pronounced peaks and troughs as illustrated in Figure 2.4 a GA seed will often lie near, but not at, a local minimum. To mitigate this, exploratory mechanisms are activated around each seed before and during local refinement: small stochastic perturbations of the solution within bounds are imposed to trigger multi-start attempts. Additional strategies such as path relinking between neighboring minima, bounded changes in segment count, and phased constraint activation increase the probability of converging to a low- Δv , constraint-satisfying minimum rather than being trapped by the initial GA seed's local topology.

3.4.1 Parallel Multi-Start with stochastic perturbations

A parallel multi start scheme with stochastic, structure aware perturbations is employed to systematically probe the neighborhood of GA generated candidates, thereby enhancing coverage of nearby basins without resorting to uninformed random restarts. This design imparts a limited “global” exploratory character to the SNOPT-based SQP refinement: because seeds already inhabit promising regions, controlled perturbations encourage the local solver to sample adjacent portions of the landscape before settling, increasing the likelihood of attaining a lower-cost feasible minimum.

In addition, this approach provide a natural avenue for parallelization by distributing independent local solves across workers, which substantially reduces wall-clock time to obtain a reduction of Δv trajectories under the imposed terminal and path constraints.

The solution vector is perturbed with the following formulation:

$$\Delta v'_i = \Delta v + u_i \quad \text{with} \quad \|u_i\| \leq \epsilon \|\Delta v_i\| \quad (3.4)$$

The parameter ϵ in figure 3.4 specifies the admissible perturbation scale; higher ϵ values induce larger deviations from the original candidate, motivating a conservative selection to ensure the trajectory remains within the solution neighborhood mapped by the GA.

3.4.2 Path relinking technique for alternative local minima identification

Path relinking is employed to expose alternative local minima by connecting an initiating GA-derived solution x^{GA} with a guiding solution x^* obtained by locally

minimizing the impulsive cost $J = \sum_{i=1}^N \|\Delta \mathbf{v}\|$ under the full set of mission constraints, and then systematically sampling and refining intermediate points along the resulting path in decision space $\Delta \mathbf{v}, \tau$.




Figure 3.4: Path relinking illustration. Credit to [13]

After identifying x^* , as it is shown in figure 3.4 a straight-line segment is constructed between x^{GA} and x^* and partitioned into n waypoints $\{x_k\}_{k=1}^n$, each of which is projected onto variable bounds and used as an initial guess for a short SQP (SNOPT). The refinement phase re-solves the same NLP, preserving both the objective and the full constraint set; by restarting from strategically varied initializations, the solver is guided to explore neighboring regions of the solution space that may yield an improved optimum.

3.4.3 Exploration of departure dates and optimal launch windows

Launch-window exploration is performed by sweeping the departure epoch t_0 and so the NRHO insertion point across multi-month spans and evaluating candidate transfers under the full constraint set to locate epochs that yield low total Δv and feasible operations. These windows are governed primarily by lunar phasing—often repeating on the lunar synodic cadence and by the chosen NRHO resonance (e.g., 9:2), which together determine when Earth–Moon geometry allows energetically favorable connections to the Gateway orbit. For direct transfers, “porkchop” analyses over t_0 and time of flight visualize cost basins and enable quick down-selection of promising dates.

For each candidate NRHO insertion epoch, a genetic algorithm is executed in backward time as the corresponding departure epoch using an identical objective and constraint set; the resulting sweep reveals when and where initiating the NRHO insertion maneuver is optimal.

3.4.4 Application of weak stability boundary (WSB) transfers

WSB transfers are incorporated as a low-energy option that exploits Sun–Earth–Moon multi-body dynamics to target weak capture at the Moon, trading longer time of flight for reduced Δv , and are therefore primarily suited for cargo-class missions within this study’s framework. In formulation, the same decision vector and objective are retained, but there is the necessity of adding a new path constraint:

- high apogee ranges beyond lunar distance
- a bound on departure energy

To realize the WSB mechanism, Sun–Earth geometry windows are enforced (e.g., constraints on the Sun–Earth–spacecraft angle or epoch ranges) while maintaining safety margins, minimum altitudes, maximum per impulse bounds, and minimum inter impulse spacing. In the global stage, the GA uses graded penalties on these geometry and energy windows alongside terminal and operational constraints, allowing near-feasible individuals that respect the WSB structure to survive selection even when minor violations persist, thereby guiding the search toward viable low-energy corridors. In the local stage, the same conditions appear as explicit inequalities in the constrained NLP solved by the SQP solver, which refines maneuver placement and segment durations to meet tight feasibility tolerances while minimizing total Δv , applying the same logic used in the case of direct transfer. Here below in figure 3.5, there is an illustration of the WSB transfer.




Figure 3.5: Weak stability boundary transfer. Credit to [2]

4. Implementation, Results, and Analysis

This chapter presents the computational implementation, the outcomes of the two stage optimization pipeline, and a focused analysis of the resulting Earth–Moon transfer trajectories. It begins by documenting the development environment and numerical setup that support high fidelity n -body propagation and constrained optimization, establishing the basis for reproducibility and performance comparisons across methods. It then reports Phase 1 (genetic algorithm) outcomes—including the behavior of single versus multiple shooting, seed selection for local refinement, and preliminary feasibility under terminal and path constraints—followed by Phase 2 (SNOPT) refinements that consolidate candidates, redistribute maneuvers, enforce constraints, and reduce total Δv .

4.1 Tool overview and logic

The tool is an integrated MATLAB pipeline for spacecraft trajectory design that consolidates high-fidelity propagation, global search, local refinement, and analysis in a coherent framework, enabling reproducible experiments across mission scenarios. Its modular architecture comprises a propagation layer (ephemerides and dynamics), an optimization layer that couples a genetic algorithm with a sequential quadratic programming solver, and a post-processing layer for metrics and visualization, all under consistent frame, unit, and time conventions. The operational logic proceeds from ephemeris-driven propagation, through GA-based diversification with penalty guidance, to SNOPT-based intensification with explicit constraints, after which diagnostics and plots are generated from the same state histories used by the solvers. This division of roles balances exploration and exploitation, improving robustness against multimodality while maintaining tight feasibility control in the final solutions. Parallel execution for multi-start and window sweeps, along with structured logging and results archiving, supports efficient search and transparent comparison across design cases.

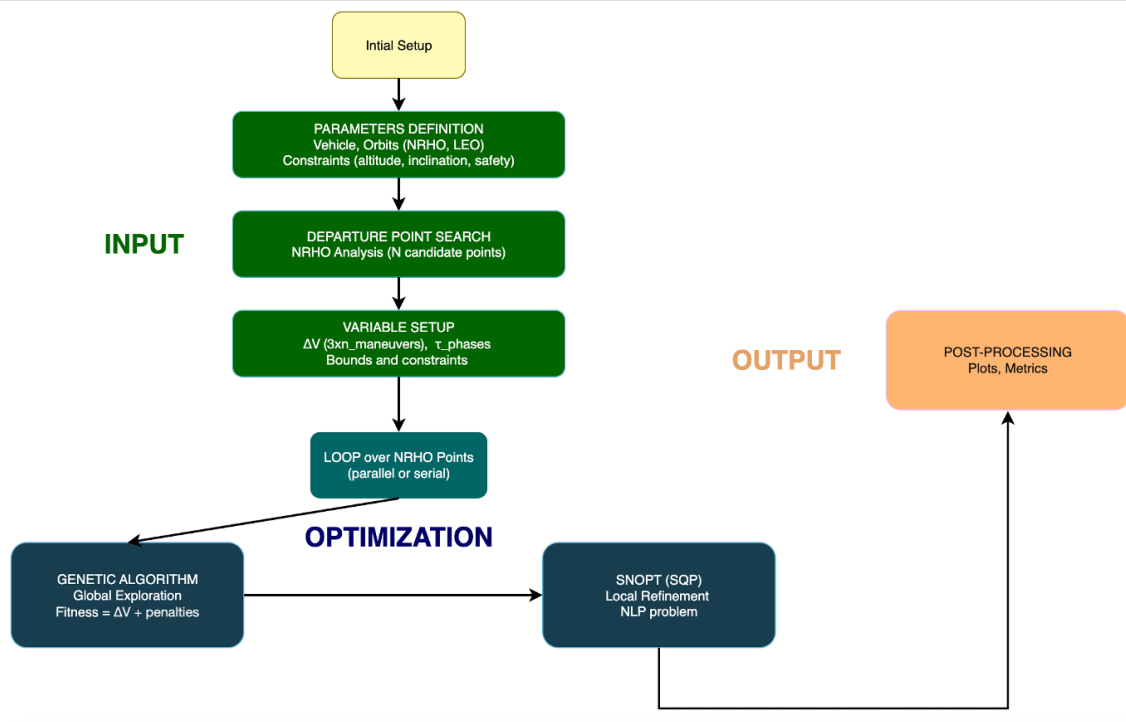


Figure 4.1: Flow chart of the the tool

4.1.1 Conventions (Frames, Time, Units, Metrics)

All trajectory propagation, multiple-shooting continuity conditions, and maneuver applications are expressed in an Earth-centered inertial frame (J2000/ICRF), ensuring that the dynamical model is free from Earth-rotation artifacts and that defect constraints are evaluated consistently. A Moon-centered inertial representation is employed when screening and annotating NRHO phases or Moon-relative distances, while Earth-fixed frames are used only for geodetic reporting (e.g., latitude/longitude/height) or ground-referenced checks. Local frames (e.g., LVLH or VNB) may be introduced for interpretive plots of maneuver geometry but are not used for propagation or optimization.

Time handling is uniform across modules: internal computations are performed on a consistent timescale (e.g. ET/TDB), which is the number of seconds since 1 January 2000. For a better human understanding, the time is converted into the standard UTC timestep. All epochs reported in figures and tables indicate the time scale used, and any conversions are applied explicitly during post-processing.

Units are standardized to kilometers (position), kilometers per second (velocity), and seconds (time); angles are reported in degrees unless otherwise indicated. Gravitational parameters, radii, and derived quantities are used consistently within this unit system. When metrics require mixed units (e.g., Δv in m/s for readability), conversions are performed only at output to maintain numerical consistency in the

solvers.

4.2 Methods shared by both cases

4.2.1 Problem formulation

The transfer is modeled as a direct, impulsive optimization in an Earth-centered inertial frame using a multiple-shooting transcription. The decision vector collects node impulses and normalized phase durations:

$$\mathbf{x} = \{\Delta \mathbf{v}_1, \dots, \Delta \mathbf{v}_N, \tau_1, \dots, \tau_{N-1}\}, \quad \Delta t_i = \tau_i T_{\text{ref}}, \quad \tau_i \in [\tau_{\min}, 1], \quad v \in [v_{\min}, v_{\max}] \quad (4.1)$$

The objective is the total impulsive cost (see eq. 3.2) subject to:

1. terminal conditions at Earth arrival (altitude and inclination ranges);
2. path/operational constraints (minimum altitude, minimum Moon-center distance, caps on per-impulse magnitude and minimum spacing, time-of-flight limit).

4.2.2 GA+SNOPT pipeline

A two-stage hybrid strategy is adopted. Stage 1 (GA) performs broad sampling with graded penalties to guide search toward near-feasible, low-cost seeds. Stage 2 (SNOPT) enforces the same constraints explicitly at tight tolerances and minimizes J (3.2) from each seed. Parallel execution is used when runs are independent; scaling and a relax-then-tighten schedule support robust convergence.

4.2.3 Diagnostics and outputs

Reported metrics include total $\sum \|\Delta \mathbf{v}\|$, time of flight, maneuver count and minimum spacing, terminal residuals, path feasibility margins, and case-specific indicators (e.g., arrival v_∞ for low-energy transfers). Visuals reuse standard plot types: 3D ECI view, lunar-frame projections, altitude profile, and maneuver timeline.

GA fitness function

Backward propagation from an NRHO state stops at Earth perigee by an ODE event detecting the ascending zero of radial velocity, $v_{\text{rad}} = \frac{\mathbf{r} \cdot \mathbf{v}}{\|\mathbf{r}\|} = 0$. The scalar fitness combines impulsive cost with penalties/rewards:

$$F(\mathbf{x}) = w_{\Delta v} \sum_{i=1}^{N-1} \|\Delta \mathbf{v}_i\| + P_{\text{alt}} + P_{\text{inc}} + P_{\text{safety}} + P_{\text{time}} + P_{\text{dir}}. \quad (4.2)$$

- Altitude window: $h_{\min}$ be the minimum geocentric altitude along the propagated arc and $[h_L, h_U]$ the target range;
- Use a piecewise penalty if $h_{\min} \notin [h_L, h_U]$, and a small negative reward toward the range center otherwise;
- Inclination range: compute $i = \arccos(h_z/\|\mathbf{h}\|)$ from $\mathbf{h} = \mathbf{r} \times \mathbf{v}$ at perigee;
- Penalize violations of $[i_L, i_U]$ and apply a modest in-range reward;
- Safety/geometry: penalize if the minimum Moon-center distance along the arc is below a buffer, if $h_{\min} < 0$ (Earth impact), or if the perigee speed exceeds local escape, $\|\mathbf{v}_{\text{peri}}\| > \sqrt{2\mu_{\oplus}/\|\mathbf{r}_{\text{peri}}\|}$;
- Phase-duration bounds: quadratic penalties on τ_i outside $[\tau_{\min}, 1]$;
- Optional regularity: light penalties to discourage cusp-like geometry, kept small relative to altitude/inclination guidance and Δv ;

Constrained NLP formulation

The local problem minimizes total impulsive cost with explicit constraints:

$$\begin{aligned} \min_{\mathbf{x}} \quad & J(\mathbf{x}) = \sum_{i=1}^{N-1} \|\Delta \mathbf{v}_i\| \\ \text{s.t.} \quad & \mathbf{c}(\mathbf{x}) \leq \mathbf{0}, \quad \mathbf{x} \in [\mathbf{l}, \mathbf{u}], \end{aligned} \tag{4.3}$$

where the constraint vector aggregates:

$$\mathbf{c}(\mathbf{x}) = \begin{bmatrix} c_{h,U} \\ c_{h,L} \\ c_{i,L} \\ c_{i,U} \\ c_{\text{Moon}} \\ c_{\oplus} \\ c_{\text{esc}} \end{bmatrix} = \begin{bmatrix} h_{\min} - h_U \\ h_L - h_{\min} \\ i_L - i \\ i - i_U \\ d_{\min, \text{Moon}, \text{req}} - d_{\min, \text{Moon}} \\ h_{\text{floor}} - h_{\min} \\ \|\mathbf{v}_{\text{peri}}\| - \sqrt{2\mu_{\oplus}/\|\mathbf{r}_{\text{peri}}\|} \end{bmatrix} \leq \mathbf{0}. \tag{4.4}$$

Bounds on τ_i and per-component $\Delta \mathbf{v}$ match the GA stage. The constraint evaluator reuses the same backward integration and perigee event to compute $h_{\min}$, i , $d_{\min, \text{Moon}}$, $\mathbf{r}_{\text{peri}}$, $\mathbf{v}_{\text{peri}}$, ensuring consistency between global guidance and local enforcement.

4.3 Direct transfer case

The first case of study is focused on the challenge of getting a crew from Earth to a Near-Rectilinear Halo Orbit (NRHO) around the Moon. For these missions, which would launch from the Guiana Space Centre in Kourou, time is of the essence. In fact, there is no possibility for astronauts to spend months in space to get there. This need for a speedy journey guided our whole approach, pushing to optimize for direct, high-thrust trajectories and leaving the more scenic, low-energy paths for uncrewed cargo flights.

4.3.1 GA results

Before initiating the detailed optimization, a preliminary analysis was conducted to identify the most promising insertion points along the NRHO. This involved a systematic scan of candidate departure points around the entire orbit, focusing on a narrow range of dates to isolate the effect of geometry. This search revealed that the optimal insertion points, in terms of minimizing propellant consumption, are located near the NRHO aposelene (approximately 70,000 km from the Moon). This is intuitively explained by the fact that the spacecraft's orbital velocity is at its lowest at this point, requiring less braking Δv to be captured into the orbit. This finding guided the selection of target points for the subsequent optimization phases.




Figure 4.2: Velocity around NRHO

The shape of this orbit led to an enormous difference between the velocity in the periselene and aposelene. It passed from $0.9 \frac{km}{s}$ to $2.65 \frac{km}{s}$. For this reason, the solver tend to apply the insertion maneuver at the aposelene.

Single shooting case

In the initially phase of the work, a single shooting method is implemented to propagate the trajectories. While simple to set up, this approach consistently resulted in solutions with high propellant consumption after a complex series of iterations.

The problem is that a single maneuver is tasked with the difficult job of correcting for both altitude and inclination at the final point of the trajectory. Achieving these two distinct and sensitive orbital parameters simultaneously with just one propulsive event is a significant challenge. In addition, this method often leads to inefficient fuel usage. Here below there is a sample figure of the result from this first part of the work. The figure 4.3 represent a lot of trajectory that started from different point from the NRHO with a single impulsive maneuver.




Figure 4.3: Trajectory from different point of the NRHO




Figure 4.4: Legend of figure 4.3

As shown in the legend of Figure 4.4, all trajectories—except for the one utilizing a lunar fly-by—are characterized by high propellant consumption, with Δv values ranging from 1000 m/s to 2000 m/s. Although these solutions respect the constraints on time of flight, they do not satisfy the stringent requirements for precision in both final altitude and inclination. In fact, as it is shown in figure 4.5 below, the black

dashed line indicates the inclination target and this constraint is never respected for each of the trajectories in figure 4.3.




Figure 4.5: Zoom of the arrival trajectories

Multiple-shooting case

For the reasons previously discussed, the difficulty of satisfying terminal constraints with a single maneuver, it became necessary to adopt a multiple shooting method for trajectory propagation. This approach provides the enhanced control fidelity required to simultaneously meet the strict requirements for both altitude and inclination at the target.

To implement this, the trajectory is discretized into N nodes within MATLAB. The propagation is achieved via a for loop, where, at the beginning of each segment, the genetic algorithm determines the three components of the velocity impulse and the propagation duration, the latter via the dimensionless time variable τ . This approach provides the algorithm with the flexibility to implicitly select the number of maneuvers, up to the maximum of $N-1$. By assigning a near-zero value to the time-scaling parameter, τ_i , the optimizer can effectively collapse a trajectory segment. This renders the maneuver at the start of that segment inconsequential, as the node is not included in the propagated path.

This approach led to an incredible decrease of Δv required for the transfer; in addition, both of the required objectives (altitude and inclination) have been achieved.




Figure 4.6: Trajectory in ECI reference system with multiple shooting method adopted

First of all, in this result it is possible to observe a consistent improvement in propellant efficiency achieved through the multiple shooting formulation. This led the solver to explore a wider part of the solution space. In fact, there is a consistent reduction in terms of required Δv that passed from $1000 \frac{m}{s}$ to $755.89 \frac{m}{s}$, representing a propellant saving of over 22%.

To visualize the solution, Figure 4.6 presents the optimized transfer trajectory (this figure and others in this section follow the same color and symbol conventions). The path is propagated backward, starting from the target NRHO, which is shown as a dashed white line for context. The journey begins at the green dot on the NRHO; this is the point where, in a real mission, the spacecraft would perform its final insertion burn. Along the way, any additional engine burns are marked with fuchsia dots.




Figure 4.7: Inclination trend w.r.t Earth during transfer




Figure 4.8: Components of velocity in RTN reference system for each maneuver

The specific mechanism by which the impulsive maneuvers alter the trajectory is shown in Figures 4.7, 4.8. These visualizations provide a deeper understanding of the control authority exerted by each maneuver and its resulting impact on the spacecraft’s flight path.

Despite being unrefined by SNOPT, the initial solutions provided by the genetic algorithm demonstrate a control history that is consistent with the principles of orbital mechanics. The optimizer correctly prioritizes maneuvers in the tangential direction, which are known to be the most efficient for modifying orbital energy. Additionally, the largest component of Δv in the normal direction—required for the plane change—is applied during the first maneuver. This corresponds to the point of apogee in the Earth-centered transfer arc, which is the dynamically optimal location to perform an inclination change. The clustering of the main propulsive events at the beginning of the transfer is a direct consequence of this strategy, as the optimizer exploits the high-altitude, low-velocity region to perform its most

significant corrections.

The clustering of the main propulsive events at the beginning of the transfer is a direct consequence of this strategy, as the optimizer exploits the high-altitude, low-velocity region to perform its most significant corrections. This new approach proved to be a major step up from the initial single-shooting attempts. Not only did it find a more fuel-efficient path, but it also successfully hit the tight landing window setted for both altitude and inclination—something the previous method couldn't do. The consequences of this are clearly visible in Figure 4.9





Figure 4.9: Keplerian orbital elements in the proximity of the Earth

4.3.2 SNOPT refinement

Following the global exploration of the solution space via the genetic algorithm, the subsequent step in our methodology is the local refinement of the most promising candidate trajectories. This section will describe the refinement techniques applied to the initial GA results and present the resulting improvements in both cost and constraint satisfaction.



Building on this result, the analysis of the components of the refined maneuvers in the local RTN (Radial-Tangential-Normal) frame was conducted, with the results shown in Figure 4.12.




Figure 4.11: Inclination trend during the transfer refined




Figure 4.12: RTN components of velocity for the refinement case

This analysis reveals that the optimizer has allocated the vast majority of the Δv to the tangential component, which is optimal for energy changes. Crucially, only a small fraction of the thrust is applied in the normal direction, which powerfully illustrates the principle that out-of-plane maneuvers are most efficient at high altitudes, where a minimal impulse can induce a change of several degrees in inclination. In

addition, the solver decides to apply only two maneuvers, with the major part of Δv concentrated on maneuver 2.




Figure 4.13: Insertion into NRHO, view in IAU MOON reference system




Figure 4.14: Keplerian orbital elements in the proximity of the Earth

A more detailed view of the trajectory’s terminal phases is provided in Figures 4.13 and 4.14, which depict the departure from the initial Earth orbit and the insertion into the NRHO, respectively.

The geometry of the final approach is shown in Figure 4.13, which represents the trajectory’s insertion into the NRHO within the IAU Moon reference frame. This figure confirms that the final maneuver is executed at the orbit’s aposelene, which is consistent with an optimal strategy to minimize the required braking Δv for capture.

4.4 Cargo Case

This case study considers an uncrewed cargo transfer from a circular Low Earth Orbit (LEO) to a Near-Rectilinear Halo Orbit (NRHO) exploiting a weak-stability-boundary (WSB) corridor to minimize total propulsive effort over an extended time of flight (TOF) of 120–150 days. The transfer is subdivided with a multiple shooting technique (see 2.3) with $N = 6$ nodes, yielding up to $N - 1 = 5$ admissible impulsive maneuvers, and segment durations parameterized by normalized variables $\tau_i \in [0.02, 1]$ scaled by a 50-day reference to promote well-conditioned timings in the optimizer. The decision vector comprises the three components of each $\Delta \mathbf{v}_i$ and the τ_i , while the objective is the minimization of total impulse (see Eq.3.2)

The constraint set encodes feasibility and WSB geometry: minimum geocentric altitude in the corridor $H \in [170, 250]$ km, arrival inclination $I \in [5^\circ, 7^\circ]$, Earth-centered apogee distance $R_{\text{apo}} \in [1.1 \times 10^6, 1.8 \times 10^6]$ km, and characteristic energy at departure $C_3 \in [-0.7, -0.3]$ km²/s². Additional safety margins enforce a minimum Earth–Moon clearance along the arc and exclude hyperbolic Earth escape at perigee. These bounds are applied in normalized form within the NLP to balance constraint scales and improve convergence.

4.4.1 GA Results

Just as it did for the crewed mission (direct transfer), our first step for the cargo mission was to figure out the cheapest way to enter the NRHO. A series of tests to see which arrival point would save the most fuel has been done. Even though this was slower using a Weak Stability Boundary trajectory, the best insertion point in terms of propellant consumption is near the perigee point, where the velocity is at its maximum in the orbit (see Fig. 4.2). This location is even more important for a cargo mission. Since the goal is to save every last drop of fuel, arriving at the fastest point will require less propellant for the insertion manoeuvre.

For this second, more complex case, the multiple-shooting method is adopted. Having already demonstrated its clear advantages in terms of both fuel efficiency and constraint satisfaction, it was the logical choice and the only method applied in this part of the analysis.

The resulting trajectory obtained with the GA solver, is represented in Figure 4.15




Figure 4.15: WBS transfer in ECI reference frame

As can be seen in Figure 4.15, the final trajectory resulting from the GA optimization resembles a Weak Stability Boundary (WSB) transfer. The spacecraft travels to an altitude of approximately 1.6 million km from the Earth centre, where it executes a low-cost plane change maneuver. Additional maneuvers are then performed to ensure the spacecraft correctly intercepts the target NRHO in the perilune region.

Figure 4.15 illustrates the key elements: the maneuvers are represented by purple dots, the Moon's orbit by a long-dashed white line, and the target NRHO by a short-dashed white line.




Figure 4.16: Components of velocity in the wbs transfer

Figure 4.16 shows the velocity-change vectors (ΔV) for each maneuver in the RTN (Radial–Transverse–Normal) frame, providing physical insight into the role of each impulse and clarifying why the optimizer selected this sequence as optimal for the case study. In addition, the figure shows that Maneuver 2 has the largest normal (N) component of the velocity change in the RTN frame, consistent with an out-of-plane plane-change action. Executing the inclination change where the orbital speed is minimal—at large radius—reduces the required plane-change, explaining why the solver places this maneuver far from Earth. Likewise, Maneuver 4 exhibits a pronounced normal (N) component for the same physical reason: at approximately 1.3×10^6 km from Earth, the lower orbital speed makes out-of-plane plane-change impulses more efficient, so the optimizer allocates the inclination correction at this large radius.




Figure 4.17: Inclination plane trend during the transfer

Figure 4.17 shows the evolution of the geocentric inclination over the entire transfer. The inclination decreases by about 30 degrees from the initial NRHO to the vicinity of Earth, with clear discontinuities at each colored marker corresponding to the applied impulsive maneuvers.




Figure 4.18: Parameter in the proximity of the Earth

In Figure 4.18, it is possible to observe the parameter in the proximity of the Earth.

4.4.2 SNOPT Refinement

Leveraging the GA-derived initial guess, the trajectory is refined with a sequential quadratic programming (SQP) solver, yielding a lower total $\Delta \mathbf{v}$. In Figure 4.19 presents the complete transfer in the Earth-Centered Inertial (ECI) frame and demonstrates a substantial reduction in total propellant consumption, i.e., a marked decrease in overall $\Delta \mathbf{v}$ achieved by the SNOPT-refined solution.




Figure 4.19: Transfer in ECI reference frame refined with SNOPT

The local SQP refinement does not substantially alter the overall trajectory geometry; rather, it delivers a significant reduction in the propellant requirement. Starting from the GA initial guess, the total $\Delta \mathbf{v}$ decreases from $219.44 \frac{m}{s}$ to $51.25 \frac{m}{s}$. As reported in the table associated with Figure 4.19, this corresponds to using only about 20% of the available propellant budget. Note that additional propellant is reserved for the Earth-return transfer.




Figure 4.20: Zoom on the insertion point in the NRHO in ECI reference frame

In Figure 4.20, the execution of the orbit-insertion maneuver is clearly illustrated. The figure resolves the burn geometry and timing that transition the spacecraft from the approach arc to the bound target orbit. By contrast with Figure 4.20, Figure 4.21 depicts the same maneuver expressed in the IAU Moon reference frame, i.e., a Moon-centered inertial frame (see Figure 3.3). The apparent geometric differences arise solely from the change of reference frame and do not reflect any alteration of the underlying dynamics.

The IAU Moon frame is translated to the lunar center and rotated relative to the inertial ECI axes, so the velocity vector is depicted with a different orientation even though the physical trajectory is unchanged. This difference is therefore due to the coordinate transformation, and is not a dynamical difference.




Figure 4.21: Insertion maneuver in IAU MOON reference frame




Figure 4.22: Components of velocity in the wbs transfer

The SQP refinement reduces the number of significant maneuvers to two (Figure 4.22), with the remaining burns serving only as small correction maneuvers. Also

in this case, the insertion maneuver—first when interpreted in backward propagation—has significant contributions in all three RTN components, with $\Delta \mathbf{v}_T$ dominant. By contrast, maneuver 3 employs only tangential-normal components ($\Delta \mathbf{v}_T$, $\Delta \mathbf{v}_N$), corresponding to a combined plane-change and apoapsis-raising burn.




Figure 4.23: Inclination plane change trend during the transfer

The inclination with respect to the Earth throughout the WSB transfer trajectory is shown in Figure 4.23. This trend is characteristic of weak stability boundary transfers. The observed oscillatory variations result from n-body perturbations as the spacecraft traverses different gravitational influence regions during this low-energy transfer. The inclination evolves from approximately 26 at departure to 7 at earth, demonstrating the inherent orbital plane reorientation that characterizes WSB transfers.




Figure 4.24: Zoom departure from Earth

4.5 Cross-Case Comparison and Analysis

The effectiveness of each trajectory optimization approach is evaluated using primary performance indicators including total propulsive cost ($\sum_{i=1}^{N-1} |\Delta \mathbf{v}_i|$ in m/s), time of flight from departure to NRHO insertion (in days), constraint satisfaction for terminal altitude and inclination requirements, and mission success rate representing the percentage of optimization runs achieving feasible solutions. Additional secondary metrics provide deeper insight into solution quality and optimization behavior: maneuver distribution capturing the number of active maneuvers and their temporal spacing, computational efficiency measured through GA generations and SNOPT iterations to convergence, robustness characterized by sensitivity to initial conditions and constraint variations, and final orbital element accuracy including semi-major axis, eccentricity, and inclination at insertion.

Table 4.1: Performance comparison between Direct and WSB transfer cases

Metric	Direct Transfer	WSB Transfer
Total Δv [m/s]	707.1	51.25
Time of Flight [days]	5.68	121.89
Number of Maneuvers	3	5
Final Altitude Error [km]	0	0
Final Inclination Error [°]	0.4	0.001
Computational Time [h]	1.2	2.5

The comparative analysis reveals distinct trade-offs between the two mission profiles. Direct transfer trajectories, designed for crewed missions, prioritize minimal time of flight at the expense of higher propellant consumption, achieving mission durations of approximately 5.68 days with a total Δv of 707.1 m/s. These trajectories typically concentrate maneuvers near the beginning of the transfer (insertion) and require stringent control over inclination and altitude at arrival. In contrast, cargo-oriented low-energy WSB transfers permit extended mission durations of 121.89 days and exploit multi-body gravitational dynamics to achieve remarkable fuel efficiency with a total Δv of only 51.25 m/s. The WSB approach achieves 92.75

The hybrid global-local optimization framework (initial genetic algorithm (GA) exploration followed by Sequential Quadratic Programming (SQP) refinement) consistently demonstrates robust convergence and the ability to navigate complex constraint landscapes. Multiple shooting, in particular, provides superior flexibility for both direct and WSB cases, enabling the optimizer to select the number and placement of maneuvers and more effectively satisfy multi-phase constraints. The effectiveness of this strategy is reflected in reduced total impulse costs and improved terminal accuracy—key metrics for both mission classes. For direct transfers, maneuvers are concentrated near departure and arrival phases with three active impulses, whereas WSB trajectories distribute five maneuvers across the extended mission duration. Both cases achieve precise terminal constraint satisfaction, with altitude errors at the target orbit effectively eliminated and inclination errors maintained within acceptable tolerances (0.4° for direct transfers, 0.001° for WSB trajectories).

For the high-energy transfers, preliminary scan analyses reveal that insertion at or near the NRHO apolune, where the spacecraft velocity is minimized at approximately 70,000 km lunar altitude, consistently yields the lowest propellant costs. Whereas, for the low-energy transfer the optimal point for the insertion maneuver is near the perilune. The selection of insertion geometry is therefore a pivotal factor in overall mission performance.

WSB-type transfers, while propellant-efficient, impose operational constraints including extended mission durations, increased sensitivity to model fidelity and timing perturbations, and higher computational burden (2.5 hours versus 1.2 hours for direct transfers). On the other hand, direct transfers offer flexibility and mission

assurance but at the cost of higher fuel consumption that impacts payload capacity and mission cost. These operational considerations must be carefully balanced against mission requirements and constraints.

5. Conclusions

The hybrid GA–SQP framework demonstrated its value in exploring the highly nonconvex trajectory design space, where the genetic algorithm efficiently explored a wide range of candidate solutions, identifying near-feasible guesses that balance propellant cost and constraint satisfaction, while SQP refinement on these guesses achieved local optimality, reducing total Δv by up to 20–30% relative to GA-only results and meeting stringent arrival altitude and inclination requirements. Multiple shooting provided the necessary flexibility, yielding robust convergence across all test cases. However, several modeling assumptions limit the current results: point-mass gravity modeling neglects higher-order lunar and terrestrial gravity harmonics, which may underestimate the total amount of Δv required; n-body dynamics omit Direct Solar Radiation Pressure (SRP) and other solar system body effects; impulsive maneuvers are assumed without modeling finite burn effects and thrust-level constraints; and fixed reference frames abstract attitude and control authority limitations of the spacecraft. To address these gaps, future work should incorporate lunar and Earth gravity field harmonics, include SRP and Jupiter gravitational accelerations in the dynamics, transition to a low-thrust or finite-burn model using direct transcription, and couple with attitude/control models to enforce realistic thrust vector limitations. Further developments should focus on three main directions: first, the inclusion of additional perturbations such as Solar Radiation Pressure (SRP) and higher-order harmonics will improve the dynamic model, enhancing prediction accuracy and giving the model reliability closer to real physics; second, the use of multi-objective formulation will allow simultaneous minimization of Δv and time of flight, or other metrics such as radiation exposure and thermal constraints, helping to establish mission requirements for future operations; third, the hybrid GA–SQP and multiple-shooting framework is flexible and can be extended to various other application cases including low-thrust transfers using electric propulsion to NRHO or Lagrange orbits, transfers to low lunar orbit (LLO), Earth–Sun Lagrange point missions such as Sun–Earth L1/L2 halo orbits, and interplanetary trajectories including Mars transfers via gravity assists.

Bibliography

- [1] Emily M. Zimovan, Kathleen C. Howell, and Diane C. Davis. “Near Rectilinear Halo Orbits and their Application in Cis-Lunar Space”. In: *3rd IAA Conference on Dynamics and Control of Space Systems (DyCoSS)*. 2017, IAA-AAS-DyCoSS3–125.
- [2] Edward A. Belbruno and John P. Carrico. *Calculation of Weak Stability Boundary Ballistic Lunar Transfer Trajectories (hosted copy)*. Hosted PDF; accessed 2025-09-17. The Astrogator’s Guild. 2000.
- [3] Gundam Italian Club. *Punti di Lagrange*. Webpage. Ultima modifica: 26 Febbraio 2021; Consultato il 11 Settembre 2025. 1979. URL: <https://www.gundamitalianclub.net/location/punti-di-lagrange/>.
- [4] Degenerate Conic. *Near-Rectilinear Halo Orbits*. Blog post. Accessed on 11 settembre 2025. 2022. URL: <https://degenerateconic.com/near-rectilinear-halo-orbits.html>.
- [5] Lorenzo Casalino. *Space trajectory optimization*. Slide del corso di Space trajectory optimization. Politecnico di Torino, 2024.
- [6] V. M. Becerra. “Practical Direct Collocation Methods for Computational Optimal Control”. In: *Modeling and Optimization in Space Engineering*. Ed. by Giorgio Fasano and János D. Pintér. New York: Springer, 2013.
- [7] Walter Ngaw. *Linear Regression*. Online blog post. Accessed: 2025-09-18. 2019. URL: <https://wngaw.github.io/linear-regression/>.
- [8] *NAIF CSPICE Toolkit Documentation*. Online documentation for CSPICE routines. NASA/JPL Navigation and Ancillary Information Facility (NAIF). 2021.
- [9] NAIF/JPL. *Frames and coordinate systems*. Technical report/tutorial. Navigation and Ancillary Information Facility (NAIF), Jet Propulsion Laboratory. Accessed September 15, 2025.
- [10] Giovanni Lavezzi, Mariusz Eivind Grøtte, and Marco Ciarcià. “Attitude Control Strategies for an Imaging CubeSat”. In: Conference paper, ResearchGate pre-print. 2019. URL: https://www.researchgate.net/publication/334308050_Attitude_Control_Strategies_for_an_Imaging_CubeSat.

-
- [11] Clark P. Newman et al. “Stationkeeping, Orbit Determination, and Attitude Control for Spacecraft in Near Rectilinear Halo Orbits”. In: *AAS/AIAA Astrodynamics Specialist Conference*. Paper AAS 18-388. Snowbird, UT, Aug. 2018.
 - [12] Advance Space. *BLT (Ballistic Lunar Transfer) Cheat Sheet*. [https://s3-us-west-2.amazonaws.com/advspace.publicshare/BLT+\(Ballistic+Lunar+Transfer\)+Cheat+Sheet.pdf](https://s3-us-west-2.amazonaws.com/advspace.publicshare/BLT+(Ballistic+Lunar+Transfer)+Cheat+Sheet.pdf). Accessed: 2025-09-11.
 - [13] ResearchGate contributor. *Path relinking illustration (Figure 1)*. Figure page; accessed 2025-09-17. ResearchGate. n.d.