

Politecnico di Torino



Master of Science Course in Aerospace Engineering

Academic year 2025-2026

MASTER OF SCIENCE THESIS

**Particle Swarm Optimisation of Multi-Impulse
Transfers in the Cislunar Environment: LEO-NRHO
and NRHO-PDI Scenarios**

Tutors

Prof. Lorenzo Casalino

Dott. Andrea Forestieri

Candidate

Philip List

Ringraziamenti

Desidero esprimere la mia sincera gratitudine al mio relatore, il Prof. Lorenzo Casalino, per avermi guidato durante tutto lo svolgimento di questo lavoro di tesi. La sua disponibilità, la competenza e i preziosi suggerimenti mi hanno consentito di affrontare con maggiore consapevolezza le scelte di modellazione e l'interpretazione dei risultati, contribuendo in modo determinante alla realizzazione di questo elaborato.

Rivolgo un sentito ringraziamento anche al mio correlatore, il Dott. Andrea Forestieri, per la disponibilità e il costante confronto durante gli incontri svolti nel corso del semestre. I suoi consigli e le sue osservazioni sono stati fondamentali nell'orientare le scelte metodologiche e computazionali adottate nel presente lavoro.

Un ringraziamento, infine, alla mia famiglia e ai miei amici, per il sostegno dimostrato lungo tutto il mio percorso di studi e per aver condiviso con me difficoltà e soddisfazioni.

*“The surface of the Earth is the shore of the cosmic ocean.
From it we have learned most of what we know.
Recently, we have waded a little out to sea,
enough to dampen our toes or, at most, wet our ankles.
The water seems inviting. The ocean calls.”*

— Carl Sagan, *Cosmos*

Abstract

This thesis addresses the optimisation of multi-impulse transfers in the cislunar environment for crewed missions, treating two complementary mission scenarios under a unified methodology: the trans-lunar leg from a Low Earth parking orbit (LEO) to the 9:2 synodic-resonant southern L_2 Near-Rectilinear Halo Orbit (NRHO), and the descent leg from the NRHO to a Powered Descent Initiation (PDI) gate above the lunar south pole. The trajectories are designed in a high-fidelity four-body dynamical model that incorporates the gravitational influence of Earth, Moon, and Sun through JPL SPICE ephemerides, and they are optimised for minimum total impulse (ΔV) by a hybrid scheme coupling Particle Swarm Optimisation (PSO) for global search with Sequential Quadratic Programming (SQP) for local refinement. Both scenarios are analysed through a cumulative-constraint methodology in which each case retains the active constraints of its predecessor and introduces a single operational requirement or additional degree of freedom, so that the marginal ΔV cost of each restriction can be isolated.

For the LEO→NRHO leg, the reference solution under the full operational constraint set converges to 3.82 km/s from Cape Canaveral and 3.84 km/s from Kourou. Releasing the convention of the insertion at the apolune is shown to enable an architectural transition from a three-impulse cislunar design to a near-two-impulse transfer, recovering 425 m/s for the Kourou launch site. A secular-variation study reveals a geometric coupling between the launch site and the transfer to the NRHO: the Cape Canaveral optimum is structurally stable across consecutive synodic periods, whereas the Kourou optimum is multi-basin and substantially more dispersed. For the NRHO→PDI leg, the south-polar descent is delivered for 0.86–0.90 km/s, with the polar-latitude constraint identified as the dominant marginal cost. The descent solutions are further found to show a structural trade-off between the arrival velocity at the PDI gate and the transfer ΔV , lower-cost transfers admitting progressively higher gate velocities.

Beyond the specific numerical values, the cumulative-constraint methodology developed here is independent of the target orbit and offers a transferable framework for quantifying the additional cost of operational requirements in future cislunar transfer design.

Sommario

Questa tesi affronta l'ottimizzazione di trasferimenti multi-impulso nell'ambiente cislunare per missioni con equipaggio, trattando due segmenti di missione complementari nell'ambito di una metodologia unificata: la tratta translunare da un'orbita di parcheggio in bassa orbita terrestre (LEO) a una Near-Rectilinear Halo Orbit (NRHO) meridionale attorno a L_2 , in risonanza sinodica 9:2, e la tratta di discesa dalla NRHO fino alla soglia di Powered Descent Initiation (PDI) al di sopra del polo sud lunare. Le traiettorie sono progettate mediante un modello dinamico a quattro corpi ad alta fedeltà, che tiene conto dell'influenza gravitazionale di Terra, Luna e Sole attraverso le effemeridi JPL SPICE, e sono ottimizzate per il minimo impulso totale (ΔV) tramite uno schema ibrido che accoppia la Particle Swarm Optimisation (PSO), impiegata per la ricerca globale, con la Sequential Quadratic Programming (SQP), impiegata per il raffinamento locale. Entrambi gli scenari sono analizzati attraverso una metodologia a vincoli cumulativi, in cui ciascun caso conserva i vincoli attivi del precedente e introduce un singolo requisito operativo o un ulteriore grado di libertà, così da isolare il costo marginale in ΔV di ciascuna restrizione.

Per la tratta LEO→NRHO, la soluzione di riferimento con l'insieme completo dei vincoli operativi converge a 3,82 km/s da Cape Canaveral e a 3,84 km/s da Kourou. Si dimostra che l'abbandono della convenzione di inserimento all'aposenio consente una transizione architetturale da tre impulsi a quasi a due impulsi, con un risparmio di 425 m/s per il sito di lancio di Kourou. Uno studio di variazione secolare evidenzia un accoppiamento geometrico tra sito di lancio e NRHO: l'ottimo di Cape Canaveral risulta strutturalmente stabile su periodi sinodici consecutivi, mentre quello di Kourou presenta bacini multipli ed è sensibilmente più disperso. Per la tratta NRHO→PDI, la discesa è ottenuta a un costo di 0,86-0,90 km/s e il vincolo di latitudine polare è identificato come il costo marginale dominante; le soluzioni di discesa descrivono inoltre un compromesso strutturale tra la velocità di arrivo alla soglia di PDI e il ΔV di trasferimento, con i trasferimenti a costo inferiore che ammettono velocità di arrivo progressivamente maggiori.

Al di là dei valori numerici specifici, la metodologia a vincoli cumulativi qui sviluppata è indipendente dall'orbita di destinazione e offre un quadro metodologico trasferibile per quantificare il costo marginale dei requisiti operativi nella progettazione di futuri trasferimenti cislunari.

Contents

1	Mission context	4
1.1	The Renewed Interest in Cislunar Space	4
1.2	The Artemis Program: past and upcoming missions	5
1.3	The Near Rectilinear Halo Orbit as an Operational Orbit	6
1.4	Human Spaceflight to the Moon: Constraints and Challenges	7
2	Astrodynamics Theoretical	
	Background	9
2.1	The N-Body Problem and High-Fidelity Dynamics	9
2.1.1	Equations of Motion in an Inertial Frame	9
2.1.2	Gravitational Model and Perturbations	11
2.1.3	Ephemerides and the JPL SPICE Toolkit	11
2.2	Orbital Environments	12
2.2.1	Low Earth Orbit and Van Allen Belts	12
2.2.2	The Circular Restricted Three-Body Problem	13
2.2.3	Lagrange Points and Halo Orbit Families	16
2.2.4	Near Rectilinear Halo Orbits: Geometry and Stability	17
2.2.5	Lunar Surface and South Pole Access and Constraints	21
3	Trajectory Optimisation:	
	Problem Formulation	23
3.1	Optimal Control Problem Statement	23
3.2	Optimisation Approaches	24
3.2.1	Thrust Model: Impulsive vs. Finite-Burn	24
3.2.2	Direct vs. Indirect Methods	25
3.2.3	Trajectory Parametrisation: the Multi-Impulse Approach	26
3.2.4	Forward and Backward Propagation	27
3.3	Optimisation Algorithms	28
3.3.1	Particle Swarm Optimisation	28
3.3.2	Sequential Quadratic Programming	30
3.3.3	Hybrid PSO-SQP Strategy	31
3.4	State of the Art	32
3.4.1	LEO–NRHO Transfer Literature	32
3.4.2	NRHO–Lunar Surface Transfer Literature	32

3.4.3	PSO in Space Trajectory Optimisation	33
4	Methodology and Problem Formulation	34
4.1	Dynamical Model Implementation	34
4.2	Reference Frames	35
4.2.1	ECI J2000 Frame	36
4.2.2	Lunar Orbital Plane–Aligned Frame	36
4.2.3	Selenocentric Inertial Frame	37
4.2.4	Body-Fixed IAU_MOON Frame	38
4.2.5	Local RTN Frame	39
4.3	Scenario 1: LEO $\rightarrow$ NRHO	40
4.3.1	Mission Definition	40
4.3.2	Decision Vector	41
4.3.3	Objective Function and Constraint Implementation	42
4.4	Scenario 2: NRHO $\rightarrow$ Powered Descent Initiation	43
4.4.1	Mission Definition	43
4.4.2	Decision Vector	45
4.4.3	Objective Function and Constraint Implementation	46
4.4.4	Extension: Free Departure Point on the NRHO	48
4.5	PSO and SQP Configuration	48
4.6	Optimisation Pipeline	52
5	Scenario 1: LEO $\rightarrow$ NRHO	54
	Roadmap of the constraint cases	54
5.1	Baseline Solution: Perigee Altitude Constraint Only	55
5.1.1	Optimal Decision Vector	56
5.1.2	Trajectory Geometry	57
5.1.3	ΔV Budget	60
5.1.4	Multi-start Convergence, Active Constraints, and Operational Look-ahead	62
5.2	Effect of LEO Inclination Constraint: Cape Canaveral (Case 3a)	65
5.2.1	Optimal Decision Vector	66
5.2.2	Trajectory Geometry	67
5.2.3	ΔV Budget: Where the Inclination Penalty Is Paid	69
5.2.4	Multi-start Convergence, Active Constraints, and Operational Look-ahead	72
5.3	Effect of LEO Inclination Constraint: Kourou (Case 3b)	74
5.3.1	Optimal Decision Vector	74
5.3.2	Trajectory Geometry	75
5.3.3	ΔV Budget and the Compound Mid-course Manoeuvre	77
5.3.4	Multi-start Convergence, Active Constraints, and Operational Look-ahead	80
5.3.5	Comparison: Cape Canaveral vs Kourou	81
5.4	Effect of Van Allen Radiation Constraint	83
5.4.1	Operational Thresholds: Cape Canaveral (Case 4a)	83
5.4.2	Operational Thresholds: Kourou (Case 4b)	85
5.4.3	Tightened Outer Belt: Cape Canaveral (Case 4c)	86
5.4.4	Tightened Outer Belt: Kourou (Case 4d)	88

5.5	Effect of the Free NRHO Arrival Point: Cape Canaveral (Case 5a)	90
5.5.1	Motivation and parametrisation	90
5.5.2	Optimal decision vector	91
5.5.3	Trajectory geometry	92
5.5.4	ΔV budget and the marginal benefit of the free arrival point	93
5.6	Sensitivity Analysis	96
5.6.1	Multi-start convergence, active constraints, and operational implications	96
5.7	Effect of the Free NRHO Arrival Point: Kourou (Case 5b)	98
5.7.1	Motivation and parametrisation	98
5.7.2	Optimal decision vector	99
5.7.3	Trajectory geometry	100
5.7.4	ΔV budget and the architectural collapse of the on-board impulse profile	101
5.8	Sensitivity Analysis	103
5.8.1	Multi-start convergence and active constraints	104
5.9	Secular Study: Cape Canaveral versus Kourou	105
5.10	Cumulative Summary	108
6	Scenario 2: NRHO to Powered Descent Initiation	111
6.1	Reference Configuration	112
6.2	Case 1: Single-Impulse Baseline	113
6.2.1	Optimal Decision Vector	113
6.2.2	Trajectory Geometry	114
6.2.3	ΔV Budget	115
6.2.4	Multi-start Convergence	115
6.3	Case 2: Three-Impulse Parametrisation, Natural TOF	116
6.3.1	Optimal Decision Vector	116
6.3.2	Trajectory Geometry	117
6.3.3	ΔV Budget	117
6.3.4	Multi-start Convergence	119
6.4	Case 3: Southern Hemisphere Arrival, Natural TOF	120
6.5	Case 4: South Polar Arrival + TOF Constraint	121
6.5.1	Optimal Decision Vector	121
6.5.2	Trajectory Geometry	122
6.5.3	ΔV Budget	122
6.5.4	Multi-start Convergence	124
6.6	Case 5: Free Departure Point on the NRHO	125
6.6.1	Three Representative Variants	125
6.6.2	Trajectory Geometries	127
6.6.3	ΔV Budget and Net-Energy Accounting	128
6.6.4	Operational Selection of the Reference Solution	129
6.6.5	Multi-start Convergence	129
6.7	Secular study	130
6.8	Cumulative Summary	131

7	Conclusions	134
7.1	Principal Results	134
7.2	Methodological Findings	135
7.3	Future Work	136
7.4	Final Remarks	137

List of Figures

1.1	Earthset	5
1.2	Cristoforetti’s return from the ISS aboard Soyuz TMA-15M, Kazakhstan	8
2.1	Geometry of a n-body system	9
2.2	Schematic visualisation of the Van Allen radiation belts (not to scale).	13
2.3	Geometry of Earth-Moon system	14
2.4	3D Zero Velocity Surfaces	15
2.5	Lagrangian points in a generic synodic RS with $\mu = 0.1$	16
2.6	Halo Orbits	17
2.7	The selected 9:2 synodic resonant NRHO in the Moon-Centred Inertial frame.	18
2.8	The selected 9:2 synodic resonant NRHO in the Earth-Centred Inertial frame.	18
2.9	The 9:2 synodic resonant southern L2 NRHO: ephemeris propagation (left) and eclipse geometry (right).	19
2.10	The stability index as a function of perilune radius across the L2 and L1 halo families, respectively.	20
3.1	Scheme of trajectory optimisation methods	23
3.2	Single shooting, multiple shooting, and the multi-impulse parametrisation used in this work	26
3.3	Distance-dependent behavioural zones in flocking models	29
4.1	ECI J2000 frame	36
4.2	Lunar orbital plane-aligned frame	37
4.3	Selenocentric inertial frame	38
4.4	Body-fixed IAU_MOON frame	39
4.5	Local RTN frame at the spacecraft state $(\mathbf{r}, \mathbf{v})$	40
4.6	PSO convergence histories for Case 4a, $N = 10$ independent seeds with deterministic initialisation $\mathbf{rng}(k)$	50
4.7	Median convergence trajectory and interquartile band	51
4.8	Detail of the last 150 iterations of each seed for the Case 4a convergence study.	51
5.1	Baseline optimal transfer in the lunar orbital plane-aligned ECI frame.	58
5.2	Close-up of the NRHO arrival region in the lunar orbital plane-aligned ECI frame.	59
5.3	Baseline optimal transfer in the selenocentric frame. The lunar sphere of influence (SOI, radius 66 000 km) is shown as a faint grid.	60
5.4	Baseline LEO departure region in the ECI J2000 frame.	62

5.5	Case 3a optimal transfer in the lunar orbital plane-aligned ECI frame, with the Cape Canaveral inclination constraint ($26^\circ \leq i \leq 31^\circ$) active.	68
5.6	Case 3a optimal transfer in the selenocentric frame.	69
5.7	Decomposition of the four impulses of Case 3a in the local Radial-Tangential-Normal (RTN) frame, computed at the pre-impulse spacecraft state.	71
5.8	Case 3a LEO departure region in the ECI J2000 frame.	72
5.9	Case 3b optimal transfer in the lunar orbital plane-aligned ECI frame, with the Kourou inclination constraint ($4^\circ \leq i \leq 8^\circ$) active.	76
5.10	Case 3b optimal transfer in the selenocentric frame.	77
5.11	Decomposition of the four impulses of Case 3b in the local Radial-Tangential-Normal (RTN) frame, computed at the pre-impulse spacecraft state.	79
5.12	Case 3b LEO departure region in the ECI J2000 frame.	80
5.13	Spacecraft altitude versus time of flight for the Case 4a optimum (Seed 6).	84
5.14	Spacecraft altitude versus time of flight for the Case 4b optimum (Seed 4).	85
5.15	Spacecraft altitude versus time of flight for the Case 4c optimum.	87
5.16	Spacecraft altitude versus time of flight for the Case 4d optimum.	89
5.17	Case 5a optimal transfer in the lunar orbital plane-aligned ECI frame, with the free NRHO arrival point active.	93
5.18	Decomposition of the four impulses of Case 5a in the local Radial-Tangential-Normal (RTN) frame, computed at the pre-impulse spacecraft state.	95
5.19	Sensitivity analysis for Case 5a	96
5.20	Case 5b optimal transfer in the lunar orbital plane-aligned ECI frame, with the free NRHO arrival point active.	101
5.21	Decomposition of the four impulses of Case 5b in the local Radial-Tangential-Normal (RTN) frame, computed at the pre-impulse spacecraft state.	103
5.22	Sensitivity analysis for Case 5b	103
5.23	Secular variation of the Case 5 free-arrival-point optima over twenty consecutive NRHO periods.	106
6.1	Case 1 trajectory, selenocentric inertial frame	114
6.2	Case 1 altitude and latitude profiles	115
6.3	Case 2 trajectory, selenocentric inertial frame	118
6.4	Case 2 altitude and latitude profiles	119
6.5	Case 4 trajectory, selenocentric inertial frame	123
6.6	Case 4 altitude and latitude profiles	124
6.7	Seed 46 trajectory	127
6.8	Seed 28 trajectory	127
6.9	Seed 5 trajectory	127
6.10	Seed 46 profiles.	127
6.11	Seed 28 profiles.	127
6.12	Seed 5 profiles.	127
6.13	Optimised departure phase α^* versus NRHO period index for the Case 5 (free- α) secular-variation study.	131

List of Tables

4.1	Decision variable bounds for Scenario 1 (LEO $\rightarrow$ NRHO). All velocity components are expressed in the lunar orbital plane-aligned frame of Section 4.2.	41
4.2	Constraint implementation for Scenario 1.	42
4.3	Decision variable bounds for Scenario 2 (NRHO $\rightarrow$ PDI). All velocity components are expressed in the selenocentric frame of Eq. (4.4).	45
4.4	Constraint implementation for Scenario 2.	46
4.5	PSO solver parameters	49
4.6	Relative variation of the PSO best fitness	50
5.1	Roadmap of the cumulative-constraint cases for the LEO $\rightarrow$ NRHO scenario.	55
5.2	Optimal decision vector $\mathbf{x}^*$ - baseline scenario (LEO $\rightarrow$ NRHO, perigee altitude constraint only).	57
5.3	ΔV budget – baseline scenario.	61
5.4	Multi-start convergence of the baseline scenario.	63
5.5	Effect of activating the Case 2 time-of-flight constraint on the multi-start optimisation outcome.	64
5.6	Sensitivity of the optimal cost to the time-of-flight bound.	65
5.7	Optimal decision vector $\mathbf{x}^*$ - Case 3a	67
5.8	ΔV budget for Case 3a, compared with the Case 1 baseline.	70
5.9	Multi-start convergence of Case 3a (Cape Canaveral inclination constraint).	73
5.10	Optimal decision vector $\mathbf{x}^*$ - Case 3b	75
5.11	ΔV budget for Case 3b, compared with the baseline (Case 1) and with the Cape Canaveral solution (Case 3a).	77
5.12	Multi-start convergence of Case 3b (Kourou inclination constraint).	81
5.13	Side-by-side comparison of the Cape Canaveral (Case 3a) and the Kourou (Case 3b) solutions	82
5.14	Comparison of the Case 3a and Case 4a optima.	84
5.15	Comparison of the Case 3b and Case 4b optima.	86
5.16	Comparison of the Case 4a (operational) and Case 4c (tightened outer-belt) optima for the Cape Canaveral branch.	88
5.17	Comparison of the Case 4b (operational) and Case 4d (tightened outer-belt) optima for the Kourou branch.	89
5.18	Optimal decision vector $\mathbf{x}^*$ for Case 5a	91
5.19	ΔV budget for Case 5a, compared with the Cape Canaveral baseline of Case 3a (apolune insertion, no Van Allen constraint) and with Case 4a (apolune insertion, Van Allen constraint active).	94
5.20	Multi-start convergence of Case 5a	97

5.21	Optimal decision vector $\mathbf{x}^*$ for Case 5b	99
5.22	ΔV budget for Case 5b, compared with the Kourou baseline of Case 3b (apolune insertion, no Van Allen constraint) and with Case 4b (apolune insertion, Van Allen constraint active).	102
5.23	Multi-start convergence of Case 5b	104
5.24	Secular variation statistics over twenty NRHO periods of the Artemis IV launch window.	107
5.25	Cumulative summary of the LEO→NRHO cases.	108
6.1	Converged decision vector for Case 1 (single-impulse baseline)	113
6.2	Multi-start convergence for Case 1: ΔV^* and PSO runtime per seed.	116
6.3	Converged decision vector for Case 2	117
6.4	Multi-start convergence for Case 2: ΔV^* and operational violation per seed.	120
6.5	Multi-start convergence for Case 3 (southern-hemisphere arrival)	121
6.6	Converged decision vector for Case 4 (south-polar arrival).	122
6.7	Multi-start convergence for Case 4 (south-polar arrival): ΔV^* and op- erational violation per seed. The lowest- ΔV seed accepted under the operational tolerance is highlighted.	125
6.8	Converged decision vectors for the three representative variants of Case 5.	126
6.9	Multi-start convergence for Case 5 (free- α).	129
6.10	Cumulative summary of the five cases analysed in Chapter 6.	131

Introduction

No human has set foot on the Moon or traveled beyond Low Earth Orbit (LEO) since the final Apollo mission (Apollo 17) in December 1972. Recently, however, there has been growing interest in returning to the Moon through initiatives such as the Artemis program. A central element of this architecture was the Lunar Gateway, a modular lunar-orbiting space station designed to support crewed missions to the Moon and eventually deeper space exploration [1]. In March 2026, NASA announced the cancellation of the Lunar Gateway in its current form, redirecting its lunar strategy toward the construction of a permanent surface base near the South Pole [2]. Nevertheless, the Near Rectilinear Halo Orbit (NRHO) remains highly relevant for future cislunar operations: its unique dynamical properties, such as long-term stability, minimal station-keeping requirements, continuous line of sight with Earth's ground stations, and regular access to the lunar surface [3], make it a natural operational orbit for any future asset deployed in the Earth-Moon system, whether crewed vehicles, logistics platforms, or staging infrastructures for deep-space missions. The relevance of NRHO as a reference orbit for human spaceflight is further confirmed by the CAPSTONE mission, which has been operating in the southern L2 NRHO since 2022 and provided meaningful contributions to the validation of autonomous navigation technologies in this orbital regime [4].

The specific NRHO of interest for this work is exactly the same adopted by the CAPSTONE mission: the southern L2 halo orbit in a 9:2 resonance with the Moon's synodic period, meaning that the spacecraft completes 9 revolutions around the Moon for every 2 lunar synodic periods. This resonance yields an orbital period of approximately 6.5 days, providing regular surface access opportunities while maintaining favourable stability properties. The orbit is characterised by a high apolune radius (70,000 km) and a low perilune radius (3,500 km from the lunar center, corresponding to an altitude of 1,760 km above the surface).

Future crewed missions operating in cislunar space will require rotating crews, since the space environment does not permit prolonged stays without significant consequences on human health, including degradation of the musculoskeletal system, cardiovascular alterations, and vision impairment. Efficient chemical propulsion is therefore essential to reduce the costs associated with the design and operations of crewed spacecraft, and to maximise the payload delivered to and from the Moon's NRHO. Furthermore, efficient transfers between NRHO and the lunar surface will allow for more sustainable and frequent access to the South Pole base.

This thesis details the methods for designing and optimising high-energy impulsive

transfers for crewed vehicles between Earth, cislunar space, and the surface of the Moon, focusing specifically on two reference scenarios: transfers between LEO and NRHO, and transfers between NRHO and the powered descent initiation (PDI) gate above the South Pole, as planned by the Artemis program. Through analysing and comparing various transfer architectures and developing optimisation techniques, the objectives of this work are structured along three main axes:

- Minimising the propulsive cost (ΔV);
- Mission-level constraints, which define the operational restrictions of the problem;
- Crew-related constraints, which define the physiological conditions that a human crew must operate within.

In the context of crewed space missions towards cislunar space planned for the coming decades, reducing the total ΔV has been chosen as the primary objective. As a matter of fact, since propellant mass is a dominant driver in spacecraft design, any reduction in ΔV has a direct impact on mission performance, enabling either a reduction in launch mass, an increase in delivered payload, or a reduction in overall mission costs.

As mentioned earlier, each one of the reference mission scenarios under investigation are defined by transfers between two well-characterised orbital environments - Earth LEO and the lunar NRHO for the first one, and the latter and the Moon's surface for the second one- from which the mission constraints are derived directly. The departure LEO is constrained by a range of altitude and inclination values, chosen to reflect realistic launch conditions consistent with launch site latitudes and current launcher performance. Instead, the target PDI above the lunar surface is defined by a latitude constraint to the South Pole region, consistent with the scientific and strategic priorities of the Artemis program. These mission constraints impose strict boundary conditions on the trajectory optimisation problem, ensuring that all solutions obtained are relevant to the practical mission applications under consideration.

In addition to ΔV minimisation and mission-level constraints, the maximum allowable time of flight is directly linked to the capacity of the human body to endure the space environment and is therefore treated as one of the most critical constraints in this work. This requirement fundamentally differentiates crewed transfer design from cargo mission design, where longer low-energy trajectories, such as Ballistic Lunar Transfers (BLT) or Weak Stability Boundary (WSB) transfers, could be exploited to significantly reduce propellant consumption. For crewed missions, the time of flight constraint typically precludes such approaches, necessitating direct high-energy transfers that balance ΔV efficiency with acceptable transit durations.

The trajectory design and optimisation are carried out using a high-fidelity four-body dynamical model, incorporating the gravitational influences of Earth, Moon, and Sun, whose orbits are propagated using planetary ephemerides from the JPL SPICE toolkit. This approach ensures that the solutions obtained are physically realistic and directly applicable to actual mission planning, without reliance on simplified models such as the Circular Restricted Three-Body Problem (CR3BP).

The complexity of the resulting optimisation problem, characterised by a large and highly sensitive solution space, highlights the need for robust numerical methods capable of identifying high-quality trajectories while avoiding convergence to local minima. To this end, a Particle Swarm Optimisation (PSO) algorithm has been implemented as the global search strategy, coupled with Sequential Quadratic Programming (SQP) as a local refinement approach. While PSO is widely regarded as one of the most promising metaheuristic techniques for space trajectory optimisation, it has not yet been applied to crewed transfers encompassing both the LEO-NRHO and NRHO-PDI segments.

The most closely related work is that of Roberto [5], who addressed low-energy LEO-NRHO transfers for the Lunar Gateway using a Genetic Algorithm (GA) combined with SQP, considering both crewed and cargo missions departing from Kourou. While sharing a similar hybrid optimisation structure, that study does not extend the analysis to the NRHO-PDI segment, which constitutes a central element of the present work. Patrick [1] further investigated fuel-optimal transfers to the NRHO using indirect optimal control methods in a high-fidelity dynamical model, focusing however on cargo resupply missions with solar electric propulsion departing from GTO. The present work therefore addresses a gap in the existing literature by applying PSO-based global optimisation to the full crewed mission sequence: a combination that, to the author’s knowledge, has not been previously investigated.

The remainder of this thesis is organised as follows. Chapter 1 introduces the overall mission context, emphasising the growing importance of cislunar space and the relevance of the NRHO as a strategic operational orbit for future human exploration. Chapter 2 provides the theoretical background and a review of the state of the art, covering the fundamental concepts of astrodynamics and the characterisation of NRHOs. Chapter 4 delves into the mathematical foundations of trajectory optimisation, including the algorithms employed. Chapter 4 describes the methodology and problem formulation, detailing the high-fidelity dynamical model, the use of ephemerides for accurate propagation, the formal definition of the optimisation problem, and the hybrid PSO-SQP strategy. Chapter 5 and Chapter 6 present respectively the implementation, results, and analysis for both the LEO-NRHO and NRHO-PDI scenarios under different constraints. Finally, Chapter 7 summarises the conclusions of this work, discusses the strengths and limitations of the approach, and outlines directions for future research.

1 | Mission context

1.1 The Renewed Interest in Cislunar Space

For more than five decades following the conclusion of the Apollo program, human space-flight was confined to Low Earth Orbit (LEO), with the International Space Station (ISS) representing the primary destination for crewed missions. As stated at the beginning of the introduction, the last time a human being ventured beyond LEO was during Apollo 17, in December 1972. This long hiatus was driven by a combination of factors: the enormous cost of deep space transportation and the absence of political will, a clear strategic direction, and the technological maturity required to sustain human presence safely in the harsh environment beyond Earth’s protective magnetosphere.

In recent years, however, a profound shift has taken place. Cislunar space - the region between the Earth and the Moon and their gravitational spheres of influence - has emerged as the primary frontier of human space exploration. Space agencies across the world, including NASA, ESA, JAXA, and CSA, have converged on a shared vision in which the Moon is not merely a destination but a proving ground and a staging point for deeper space exploration, ultimately including crewed missions to Mars. The lunar surface, particularly the South Pole, is of exceptional scientific and strategic interest due to the confirmed presence of water ice in permanently shadowed craters, which could be exploited as a resource for life support systems and as a propellant feedstock via in-situ resource utilization (ISRU). This convergence is further reinforced by the growing involvement of commercial actors and by an intensifying geopolitical competition - most notably between the United States and China through the Artemis and Chang’e programmes respectively - which together accelerate the pace of cislunar mission development [6].

From a trajectory design perspective, cislunar space presents fundamentally different challenges compared to LEO operations. The dynamics of the Earth-Moon system are governed by the interplay of multiple gravitational bodies, making the problem highly nonlinear and sensitive to initial conditions. Unlike LEO, where two-body dynamics provide a good first approximation, missions in cislunar space require the use of multi-body dynamical models to accurately capture the spacecraft’s motion and to identify energetically favorable transfer opportunities [7]. This complexity, combined with the critical nature of crewed missions, makes trajectory optimisation an essential discipline in any cislunar mission architecture, with growing relevance as both governmental and commercial actors increase their operational presence in this region.

1.2 The Artemis Program: past and upcoming missions

Despite the commercial sector being a crucial player in the development of future cislunar missions, the Artemis program, led by NASA in collaboration with ESA, JAXA, CSA, and a number of commercial partners, represents the primary institutional framework for the return of humans to the Moon [8]. The program is structured as a progressive sequence of missions of increasing complexity. Artemis I, launched in November 2022, was an uncrewed test flight of the Space Launch System (SLS) and the Orion spacecraft around the Moon, completed successfully. Artemis II (April 2026) was the first crewed mission in which four astronauts completed a free-return trajectory around the Moon, reaching a record distance of approximately 407,000 km from Earth before splashing down after ten days [9]. Artemis III, planned for mid-2027, will be a crewed mission in Low Earth Orbit designed to validate rendezvous and docking operations between the Orion spacecraft and one or both commercial Human Landing Systems - SpaceX Starship HLS and Blue Origin Blue Moon - as well as the Axiom Space AxEMU extravehicular suit [2]. Artemis IV, targeted for May 2028, will be the first crewed lunar landing of the program, with a descent to the South Pole region.



Figure 1.1. "Earthset" captured from the Orion spacecraft window during the Artemis II lunar flyby, on April 6, 2026 (Flight Day 6). Credit: NASA

The original Artemis architecture envisioned the Lunar Gateway - a modular, crewed space station in a Near Rectilinear Halo Orbit (NRHO) around the Earth-Moon L2 Lagrange point - as a central node for lunar surface operations [10, 11]. The Gateway was conceived to serve as a staging platform for lunar landers, a laboratory for long-duration deep-space experiments, and a logistics hub for future missions to Mars. However, in March 2026, NASA announced the cancellation of the Lunar Gateway, citing cost overruns,

schedule delays, and the strategic decision to prioritise direct access to the lunar surface. NASA’s revised plan allocates approximately \$20 billion over seven years to build a permanent base near the lunar South Pole, progressively developed across multiple missions and ultimately supporting rotating crews on a semi-continuous basis.

Despite the cancellation of the Gateway, the NRHO retains its strategic and operational relevance for cislunar missions. The dynamical properties of this orbit make it a natural operational location for any future cislunar asset, independently of the specific programmatic context [3, 12]. Moreover, NASA’s stated goal of conducting regular crewed landing missions from 2028 onward places an even greater importance on the availability of optimised, repeatable, and operationally robust transfer trajectories between Earth and the lunar vicinity, regardless of whether an orbiting waystation is present.

1.3 The Near Rectilinear Halo Orbit as an Operational Orbit

Near Rectilinear Halo Orbits (NRHOs) are a family of periodic orbits in the vicinity of the Earth–Moon collinear Lagrange points, characterised by a highly elongated, nearly rectilinear geometry [13]. The spacecraft spends the majority of its orbital period far from the Moon at high apolune altitudes, and makes brief, close passes near the lunar poles at low perilune altitudes. This geometry, combined with the gravitational dynamics of the Earth–Moon–Sun system, gives NRHOs a set of properties that are uniquely favourable for cislunar operations [11].

The NRHO selected as the reference orbit for this work is the southern L2 halo orbit in a 9:2 resonance with the Moon’s synodic period: the spacecraft completes 9 revolutions around the Moon for every 2 revolutions that the Moon completes around the Earth, yielding an orbital period of approximately 6.5 days [14]. This resonance structure provides long-term orbital stability with very low station-keeping ΔV costs, a continuous or near-continuous line of sight with Earth’s ground stations, and regular, predictable access windows to the lunar surface [15]. These properties are intrinsic to the orbit’s dynamics and make the NRHO one of the reference orbits for any future crewed or uncrewed asset deployed in cislunar space.

The operational viability of the NRHO has been validated in flight by the CAPSTONE mission (Cislunar Autonomous Positioning System Technology Operations and Navigation Experiment), a 12U CubeSat developed by Advanced Space under a NASA contract [4, 16]. Launched on 28 June 2022 and inserted into the southern L2 NRHO on 13 November 2022, CAPSTONE has demonstrated semi-autonomous orbital maintenance, validated navigation and station-keeping technologies, and confirmed that the dynamical stability predictions for the 9:2 NRHO are consistent with reality [17]. This mission provides direct empirical evidence that the NRHO is a practical and manageable operational environment.

1.4 Human Spaceflight to the Moon: Constraints and Challenges

The design of crewed transfer trajectories in cislunar space is subject to constraints that have no direct analogue in robotic mission design. The human body is not well-suited to prolonged exposure to the space environment, and the physiological consequences of deep-space missions impose strict operational limits that directly shape the trajectory design problem.

The primary physiological effects of relevance for transit missions include:

- degradation of the musculoskeletal system due to the absence of gravitational loading in microgravity, with bone density losses estimated at approximately 2.3% per month reported for long-duration ISS missions [18];
- cardiovascular deconditioning, including cephalad fluid shifts and reduction in cardiac output;
- visual impairment associated with elevated intracranial pressure, a condition known as Spaceflight-Associated Neuro-ocular Syndrome (SANS), whose clinical manifestations include optic disc oedema, globe flattening, and hyperopic refractive shifts [19, 20].

In addition, exposure to galactic cosmic radiation beyond the protection of the Earth’s magnetosphere represents one of the most significant health risks for cislunar and deep-space missions, and is known to exacerbate bone loss beyond the effects of microgravity alone [21]. These effects collectively impose a maximum acceptable time of flight, which in this work is treated as a hard constraint on the trajectory optimisation problem.

Regarding the impact of radiations, in addition to the constraints derived from physiological effects, a further time constraint is imposed on the trajectory with respect to the transit through the Van Allen radiation belts. These are two toroidal regions of highly energetic charged particles — primarily protons in the inner belt and electrons in the outer belt - trapped by Earth’s geomagnetic field at altitudes ranging from approximately 1000 to 6,300 km (inner belt) and from 10,000 to 40,000 km (outer belt). The total ionising dose accumulated by a crew is directly proportional to the time spent within these regions: while a fast, direct transfer trajectory results in transit times on the order of hours and negligible crew exposure, low-thrust spiral trajectories that spend weeks to months within the belts can accumulate doses orders of magnitude higher [22]. For this reason, a maximum allowable transit time through the Van Allen belts is included as an additional hard constraint in the trajectory optimisation problem.

This time-of-flight constraint has a fundamental consequence for the transfer architecture: it rules out the use of low-energy transfer families, such as Ballistic Lunar Transfers (BLT) and Weak Stability Boundary (WSB) trajectories, which typically require flight times of 60–100 days in exchange for significant ΔV savings. While such trajectories are well-suited for cargo and resupply missions, they are incompatible with the physiological limits of a human crew. Crewed transfers must therefore rely on high-energy impulsive manoeuvres, accepting a higher propellant cost in exchange for transit times that are



Figure 1.2. ESA astronaut Samantha Cristoforetti, member of the International Space Station’s Expedition 43 crew, receives assistance while exiting the Soyuz TMA-15M capsule shortly after touchdown in a remote Kazakh landing zone on 11 June 2015, marking the end of her long-duration mission aboard the station. ¹

compatible with crew health and safety. The minimisation of ΔV within this constrained framework is a non-trivial optimisation problem: the solution space is large and complex, the dynamics are highly sensitive, and the constraints are numerous and heterogeneous. Therefore, the core challenge addressed in this thesis lies in balancing propellant efficiency, operational feasibility, and crew safety.

At the same time, the NRHO itself is a dynamically complex environment: the timing of insertion and departure from the orbit may significantly affect the delta-V required for the manoeuvres [11, 15]. Together, these considerations define a highly constrained, multi-scenario optimisation problem that demands advanced numerical methods and a robust global search strategy, motivating the methodological choices described in the following chapters.

¹Source: <https://www.nasa.gov/humans-in-space/walk-the-line-nasa-studies-physical-performance-after-spaceflight/>

2 | Astrodynamics Theoretical Background

This chapter establishes the astrodynamical foundations required to formulate the trajectory design problems addressed in this thesis. The discussion progresses from the equations of motion governing spacecraft dynamics in a multi-body gravitational field, to the high-fidelity four-body ephemeris model adopted for propagation, to the characterisation of the orbital environments directly relevant to the two mission scenarios.

2.1 The N-Body Problem and High-Fidelity Dynamics

2.1.1 Equations of Motion in an Inertial Frame




Figure 2.1. Geometry of a n-body system

The motion of a spacecraft subject to the gravitational attraction of n massive bodies is governed by the n -body problem. The equations of motion of the spacecraft, expressed

relative to the primary body P_i in an Earth-Centered Inertial (ECI) frame, take the form:

$$\ddot{\mathbf{r}}_{iq} = -G \frac{(m_i + m_q)}{r_{iq}^3} \mathbf{r}_{iq} - G \sum_{\substack{j=1 \\ j \neq i, q}}^N m_j \left(\frac{\mathbf{r}_{jq}}{r_{jq}^3} + \frac{\mathbf{r}_{ij}}{r_{ij}^3} \right) \quad (2.1)$$

where G denotes the universal gravitational constant, $m_{(\cdot)}$ is the mass of the body $P_{(\cdot)}$, $\mathbf{r}_{iq}$ denotes the position of the spacecraft relative to the primary body, $\mathbf{r}_{jq}$ the position of the spacecraft relative to a perturbing body, and $\mathbf{r}_{ij}$ the position of a perturbing body relative to the primary, consistent with Figure 2.1. Note, a bold letter indicates a vector entity, while an identical non-bold symbol refers to the modulus of the corresponding vector. The right-hand side of Equation (2.1) consists of two distinct terms, each with a clear physical meaning, as detailed below. The first term,

$$-G \frac{(m_i + m_q)}{r_{iq}^3} \mathbf{r}_{iq},$$

represents the two-body gravitational interaction between the spacecraft m_q and the primary body m_i (i.e., the Earth in the ECI frame). The acceleration is directed opposite to $\mathbf{r}_{iq}$ and thus oriented toward m_i , as indicated by the negative sign. The second term,

$$-G \sum_{\substack{j=1 \\ j \neq i, q}}^N m_j \left(\frac{\mathbf{r}_{jq}}{r_{jq}^3} + \frac{\mathbf{r}_{ij}}{r_{ij}^3} \right),$$

accounts for the gravitational perturbations exerted by all remaining $N - 2$ bodies (e.g., the Moon, the Sun, or other planets). The term $\mathbf{r}_{jq}/r_{jq}^3$ represents the direct gravitational acceleration of m_j on the spacecraft m_q , while $\mathbf{r}_{ij}/r_{ij}^3$ is the acceleration of m_j on the primary body m_i , which represents the indirect effect due to the acceleration of the reference frame itself. Both contributions appear here with the same overall negative sign, since $\mathbf{r}_{jq}$ is now defined as the spacecraft's position relative to the perturbing body. For a distant perturber ($r_{jq} \gg r_{iq}$), the vectors $\mathbf{r}_{jq}$ and $\mathbf{r}_{ij}$ become nearly antiparallel and comparable in magnitude, so their sum nearly vanishes: the indirect (frame) effect very nearly cancels the direct effect, and the net perturbative contribution of a distant body m_j is correspondingly reduced.

Assuming the spacecraft mass is negligible with respect to the primary bodies (restricted problem), and taking the Earth as the central body ($P_i \equiv P_{\oplus}$), Eq. (2.1) becomes:

$$\ddot{\mathbf{r}}_{\oplus, \text{sc}} = -\frac{\mu_{\oplus}}{r_{\oplus, \text{sc}}^3} \mathbf{r}_{\oplus, \text{sc}} - \sum_{\substack{j=1 \\ j \neq \oplus, \text{sc}}}^N \mu_j \left(\frac{\mathbf{r}_{j, \text{sc}}}{r_{j, \text{sc}}^3} + \frac{\mathbf{r}_{\oplus, j}}{r_{\oplus, j}^3} \right) \quad (2.2)$$

where $\mathbf{r}_{\oplus, \text{sc}}$ is the spacecraft position vector relative to Earth's centre, $\mu_{\oplus} = GM_{\oplus}$ is Earth's gravitational parameter, $\mu_j = GM_j$ and $\mathbf{r}_{\oplus, j}$ are the gravitational parameter and position vector of the j -th perturbing body relative to Earth respectively, $\mathbf{r}_{j, \text{sc}}$ is the spacecraft position vector relative to the j -th perturbing body, and $r = |\mathbf{r}|$ denotes the Euclidean

norm¹. The first term on the right-hand side represents the gravitational acceleration due to Earth; the second term accounts for the direct attraction of each perturbing body j on the spacecraft, $\mu_j \mathbf{r}_{j,sc}/r_{j,sc}^3$, together with the indirect term $\mu_j \mathbf{r}_{\oplus,j}/r_{\oplus,j}^3$ that accounts for the non-inertial acceleration of the Earth itself due to the same body. Both terms carry the same overall negative sign, consistent with Eq. (2.1); for a distant perturber they become nearly antiparallel and comparable in magnitude, so their combined contribution to the spacecraft’s acceleration is strongly attenuated.

In this work, the set of perturbing bodies comprises the Moon ($j = 2$) and the Sun ($j = 3$), yielding a four-body dynamical model. Higher-order gravitational harmonics, solar radiation pressure, and planetary perturbations beyond the Sun are not included. The rationale for this choice is discussed in next section.

2.1.2 Gravitational Model and Perturbations

The dominant dynamical effect in cislunar transfers is the gravitational attraction of the Moon, which cannot be treated as a small perturbation on a geocentric orbit once the spacecraft approaches the lunar sphere of influence (radius approximately 66 200 km)².

The inclusion of the Sun is necessary to correctly capture the effects that modify the energy of the orbit and that govern the geometry of transfers involving large apogee excursions [13].

Additional perturbations - lunar and terrestrial gravity harmonics, solar radiation pressure, and the gravitational influence of Jupiter - are not modelled in this work. For the transfer time scales considered (days), these effects produce corrections of the order of a few tens of metres per second or less on the total ΔV budget [11], and their omission is consistent with the level of fidelity adopted in preliminary mission analysis. The present model nonetheless constitutes a significant improvement over the Circular Restricted Three-Body Problem (CR3BP), which is discussed in Section 2.2.2, by incorporating time-varying positions based on ephemeris for all primary bodies and by capturing Sun-Earth-Moon geometry across the full transfer duration.

2.1.3 Ephemerides and the JPL SPICE Toolkit

The positions and velocities of the primary bodies are computed at each integration step using the SPICE toolkit developed by NASA’s Navigation and Ancillary Information Facility (NAIF) [23, 24]. SPICE provides a consistent framework for handling time systems, reference frame transformations, and body state retrieval from pre-computed ephemeris kernels. In this work, the following kernel types are employed:

¹The Euclidean norm of a vector $\mathbf{r} = (x, y, z)^\top \in \mathbb{R}^3$ is defined as $|\mathbf{r}| = \sqrt{x^2 + y^2 + z^2}$, and represents the straight-line distance from the origin of the reference frame to the point (x, y, z) .

²The lunar SOI radius is computed via the Laplace approximation $r_{\text{SOI Moon}} = r_{\text{Moon-Earth}} \left(\frac{\mu_{\text{Moon}}}{\mu_{\oplus}} \right)^{2/5}$, which assumes a spherical influence boundary.

- **SPK kernels:** `de440.bsp` provides position and velocity of the Moon and Sun; `receding_horiz_3189_1burnApo_DiffCorr_15yr.bsp` provides the reference NRHO trajectory states;
- **LSK kernel:** `naif0012.tls` enables consistent conversion between UTC and ephemeris time (ET);
- **PCK kernel:** `pck00011.tpc` is required to construct the body-fixed IAU_MOON frame, used to evaluate the selenographic latitude constraint at the PDI gate.

All trajectory propagation and constraint evaluations are performed in the J2000 ECI frame. State transformations to Moon-centred frame are applied only at post-processing stage and are never introduced into the integrator’s right-hand side, which avoids per-step coordinate rotations and the associated computational cost.

2.2 Orbital Environments

2.2.1 Low Earth Orbit and Van Allen Belts

Low Earth Orbit (LEO) is conventionally defined as the orbital regime between approximately 160 km and 2000 km altitude above Earth’s surface. In this work, the departure orbit for Scenario 1 (LEO $\rightarrow$ NRHO) is a circular orbit at an altitude between 160 km and 600 km. Two launch sites are considered: Kennedy Space Center (KSC) at Cape Canaveral in Florida ($\delta \approx 28.5^\circ$), and the Guiana Space Centre (CSG) at Kourou in French Guiana ($\delta \approx 5.2^\circ$), yielding inclination ranges of 26° – 31° and 4° – 8° respectively [25]. In this regard, it is worth mentioning that the orbital inclination achievable from a given launch site is constrained by its latitude δ through the launch geometry relation $\cos(i) = \cos(\delta) \sin(\psi)$, where ψ is the launch azimuth. Without post-insertion manoeuvres, the minimum achievable inclination equals the site latitude ($i_{\min} = \delta$), obtained for a due-east launch ($\psi = 90^\circ$). Small inclination reductions below δ are nonetheless achievable via plane-change manoeuvres once in orbit; the associated ΔV penalty remains modest and may be outweighed by the savings on the subsequent transfer to the Moon, whose orbital plane imposes its own alignment requirements on the departure geometry. Higher inclinations are accessible by rotating the azimuth, at the cost of reduced benefit from Earth’s rotation.

The choice of the 160-600 km altitude range is not arbitrary either: it is directly linked to the presence of the Van Allen radiation belts, which impose strict constraints on both the staging orbit and the transfer trajectory for crewed missions.




Figure 2.2. Schematic visualisation of the Van Allen radiation belts (not to scale).

The Van Allen radiation belts are two toroidal regions of highly energetic charged particles trapped by Earth’s geomagnetic field [26]. As shown in Figure 2.2, the inner belt, extending from approximately 1000 km to 6300 km altitude, is dominated by high-energy protons with energies up to hundreds of MeV; the outer belt, spanning from approximately 10000 km to 40000 km, consists primarily of electrons. The two belts are separated by the slot region (~6300–10000 km), characterised by significantly lower particle flux.

The total ionising dose accumulated by a crew member is directly proportional to the time spent within these regions [26].

This has two direct implications for the trajectory design. First, the upper bound of 600 km for the departure orbit altitude is set below the onset of the inner belt at ~1000 km: a crewed staging orbit above this threshold would expose the crew to significant radiation prior to Trans-Lunar Injection, which is operationally unacceptable. Second, the transfer trajectory itself must transit through the Van Allen belts during departure; the time spent within the inner belt region (1000–6300 km) is therefore imposed as a hard constraint on the optimisation problem. Both belts are taken into account for total crew exposure, albeit the inner belt poses a significantly greater biological hazard than the outer belt, whose electron population deposits a comparatively lower dose in human tissue. The transit-time constraint is therefore applied separately to the inner belt region and to the outer region.

The choice of a circular departure orbit is motivated by operational considerations: it minimises the sensitivity of the total ΔV to the in-plane argument of departure, and is consistent with standard practice for human spaceflight mission analysis [25].

2.2.2 The Circular Restricted Three-Body Problem

Although the dynamical model adopted for trajectory optimisation in this thesis is a full four-body ephemeris model, the Circular Restricted Three-Body Problem (CR3BP) constitutes the theoretical foundation on which the classification of cislunar orbit families

- including the NRHO - is based, and on which the majority of the relevant literature relies. It is therefore presented here both for completeness on the matter of NRHOs and to contextualise the transition to the higher-fidelity model.




Figure 2.3. Geometry of Earth-Moon system

The CR3BP models the motion of a massless spacecraft under the gravitational attraction of two primary bodies (here, Earth and Moon) that revolve on circular orbits about their common barycentre, as illustrated in Figure 2.3 from [13]. Expressed in the dimensionless synodic (co-rotating) reference frame, the equations of motion take the form:

$$\begin{cases} \ddot{x} - 2\dot{y} = \frac{\partial U}{\partial x} \\ \ddot{y} + 2\dot{x} = \frac{\partial U}{\partial y} \\ \ddot{z} = \frac{\partial U}{\partial z} \end{cases} \quad (2.3)$$

where the pseudo-potential U is defined as:

$$U = \frac{1}{2} (x^2 + y^2) + \frac{1 - \mu}{r_1} + \frac{\mu}{r_2} \quad (2.4)$$

with $\mu = M_{\text{Moon}} / (M_{\oplus} + M_{\text{Moon}}) \approx 0.01215$ the mass ratio of the Earth-Moon system, and r_1, r_2 the distances from the spacecraft to Earth and Moon respectively. The terms $-2\dot{y}$ and $2\dot{x}$ are the Coriolis accelerations arising from the rotating frame, while the $x^2 + y^2$ term in U accounts for the centrifugal contribution. Since the equations of motion are

autonomous³, they admit the Jacobi integral as a conserved quantity:

$$C_J = 2U - v^2 \quad (2.5)$$

where v is the spacecraft speed in the synodic frame. The Jacobi constant defines the boundaries of the regions of space that the spacecraft can physically access: for a given value of C_J , the spacecraft is confined to the region where $2U \geq C_J$, i.e. where the kinetic energy is non-negative. The surfaces where $v = 0$, known as *zero-velocity surfaces*, mark the boundary of this accessible region, as shown in Figure 2.4, taken from [27].




Figure 2.4. 3D Zero Velocity Surfaces

The CR3BP is not employed for trajectory propagation in this work, for two reasons. First, its time-invariance makes it unable to capture the secular effects of solar gravity (as anticipated in section 2.1.2) and the true geometry of the Earth-Moon-Sun system, which are relevant for transfers spanning several days. Second, and more fundamentally, the use of ephemeris-based initial conditions for the NRHO requires a dynamical model that is consistent with real-world epochs, which the autonomous CR3BP cannot provide. Nonetheless, the CR3BP remains the reference framework for understanding the orbit families, stability properties, and resonance structure of the NRHO, as described in the following sections.

³A system of differential equations is said to be autonomous if time does not appear explicitly as an independent variable, i.e. it takes the form $\dot{\mathbf{x}} = f(\mathbf{x})$ rather than $\dot{\mathbf{x}} = f(\mathbf{x}, t)$. In the synodic frame, the primaries are fixed by construction, so the equations of motion depend only on the spacecraft state $\mathbf{x}$, not on the epoch.

2.2.3 Lagrange Points and Halo Orbit Families




Figure 2.5. Lagrangian points in a generic synodic RS with $\mu = 0.1$

In the CR3BP, five equilibrium points exist at which the spacecraft experiences zero net force in the synodic frame. These are the Lagrange (or libration) points L_1 – L_5 , with L_1 , L_2 , and L_3 lying on the Earth–Moon line (collinear points) and L_4 , L_5 forming equilateral triangles with the two primaries, as shown in Figure 2.5, taken from [27]. In the Earth–Moon system, L_2 is located approximately 61 500 km beyond the Moon along the Earth–Moon line, at a distance that allows a spacecraft to maintain an approximately continuous line of sight with both the Earth and the lunar far side.

Around the collinear points, families of periodic and quasi-periodic orbits exist that have received extensive attention in mission design. The principal families include planar Lyapunov orbits, three-dimensional halo orbits, and their quasi-periodic generalisations - Lissajous orbits [13]. The halo orbit concept was first introduced by Farquhar in 1968 to describe a three-dimensional periodic orbit that maintains a constant line of sight to the Earth and is never obscured by the Moon [13]. The L_2 halo family was later formally continued by Breakwell and Brown (1979), who demonstrated the existence of a connected family of solutions evolving from the planar Lyapunov orbit near the respective Lagrange point all the way to the near-vicinity of the Moon [13].

Halo orbits bifurcate from the Lyapunov family at a specific value of the Jacobi constant when an out-of-plane displacement is introduced; they are therefore members of a three-dimensional branch that does not exist in the planar problem. As orbit size increases and perilune altitude decreases, the geometry transitions progressively from near-planar to nearly rectilinear, with the orbit becoming increasingly elongated and out-of-plane [13], as illustrated in Figure 2.6, taken from [13]. The subsection of the halo family that lies closest to the smaller primary and exhibits this characteristic elongated geometry is the Near Rectilinear Halo Orbit (NRHO) family.




Figure 2.6. Halo Orbits around L_1 and L_2

2.2.4 Near Rectilinear Halo Orbits: Geometry and Stability

Near Rectilinear Halo Orbits (NRHOs) occupy the extreme end of the L_1 and L_2 halo families, representing the subset closest to the primary whose linear stability characteristics are formally bounded by a sequence of bifurcations across the halo family [13, 28]. As anticipated, within these bounds, the orbital geometry transitions to a highly elongated, nearly polar shape: the spacecraft spends the majority of its orbital period far from the Moon at high apolune altitudes, and makes brief, rapid passes near the lunar poles at low perilune altitudes [13, 11]. It is this combination of extreme eccentricity and out-of-plane orientation that gives the NRHO its characteristic “near-rectilinear” appearance in the rotating frame.

A further operational advantage of synodic resonance arises from the orbit’s geometry in the Sun–Moon rotating frame. For a spacecraft in a $y:k$ synodic resonant NRHO completing y revolutions per k lunar synodic periods, the trajectory forms a pattern of y lobes in the Sun–Moon rotating frame, with the lunar shadow fixed along the positive x -axis [13]. By selecting the insertion epoch such that the shadow aligns with the gaps between lobes, lunar eclipse events can be completely eliminated. Orbits with low $y:1$ resonances, such as the 3:1 and 4:1 synodic resonant NRHOs, are particularly effective for eclipse avoidance due to their large gaps between lobes, while the 9:2 resonance offers a compromise between eclipse management and perilune altitude [13, 28].

The reference orbit for this work is the southern L_2 halo orbit in a 9:2 synodic resonance with the Moon’s period. The spacecraft completes 9 revolutions over 2 lunar synodic periods (≈ 59.06 days), yielding an individual orbital period of approximately 6.5 days [14]. This orbit is characterised by a perilune radius of approximately 3 158 km and an apolune radius of approximately 71 000 km [13, 11], as shown in Figure 2.7. In this work, the initial conditions for the NRHO are obtained directly from ephemeris data via SPICE using the kernel `receding_horiz_3189_1burnApo_DiffCorr_15yr.bsp`, rather than from a CR3BP model, ensuring consistency with real mission epochs [23]. The same orbit, expressed in the Earth-Centred Inertial frame adopted for trajectory propagation throughout this work, is shown in Figure 2.8.




Figure 2.7. The selected 9:2 synodic resonant NRHO in the Moon-Centred Inertial frame.





Figure 2.8. The selected 9:2 synodic resonant NRHO in the Earth-Centred Inertial frame.


Eclipse Properties

A distinctive operational advantage of the 9:2 synodic resonant NRHO is its favourable eclipse geometry. Due to the near-rectilinear shape of the orbit, shown in Figure 2.9a as propagated under the ephemeris model adopted in this work, the spacecraft spends the vast majority of each revolution at high altitude above the south lunar pole, far from the lunar shadow cone. Eclipse events, when they occur, are limited to approximately 60 min per 6.5-day revolution, representing a small fraction of the total orbital period [13].

Moreover, eclipse periods can be entirely avoided by selecting the orbital insertion epoch such that the shadow cone aligns with the gaps between successive orbital lobes, as illustrated in Figure 2.9b. This property is particularly relevant for solar-powered spacecraft and crewed vehicles, for which prolonged loss of solar power poses an operational constraint.



(a) 9:2 NRHO propagated in the Moon-centred inertial frame using the ephemeris model adopted in this work.



(b) 9:2 NRHO in the Moon-centred inertial frame showing the lunar shadow cone [13].

Figure 2.9. The 9:2 synodic resonant southern L2 NRHO: ephemeris propagation (left) and eclipse geometry (right).

Stability Index

Within the context of the CR3BP, the linear stability of a periodic orbit is assessed by analysing how small perturbations evolve over one complete revolution. This is formalised through the State Transition Matrix (STM), a linear map that propagates deviations from a reference trajectory over a given time interval. When evaluated over precisely one orbital period, the STM is known as the monodromy matrix [13, 28]: its eigenvalues encode whether perturbations grow, decay, or remain constant in magnitude after each revolution.

Two structural properties of the monodromy matrix are relevant to the stability analysis [28]:

1. eigenvalues always appear in reciprocal pairs $(\lambda_i, 1/\lambda_i)$, a consequence of the Hamiltonian structure of the CR3BP ⁴;

⁴In a Hamiltonian system, the flow preserves volume in phase space (Liouville's theorem). This

2. one eigenvalue pair is always equal to unity, associated with the periodicity of the orbit. This trivial pair carries no stability information and is excluded from the analysis.

The remaining two reciprocal pairs are each condensed into a single scalar, the stability index [28]:

$$\nu_i = \frac{1}{2} \left(\lambda_i + \frac{1}{\lambda_i} \right), \quad i = 1, 2. \quad (2.6)$$

By construction, $|\nu_i| = 1$ when $|\lambda_i| = 1$, i.e. when the perturbation neither grows nor decays. An orbit is therefore linearly stable if and only if both indices satisfy $|\nu_i| \leq 1$; otherwise it is unstable, with larger values of $|\nu_i|$ indicating faster departure from the reference path.




Figure 2.10. The stability index as a function of perilune radius across the L2 and L1 halo families, respectively.

Halo orbits near the Lagrange points exhibit stability indices significantly greater than one, indicating strongly unstable behaviour with divergence on sub-revolution timescales. As perilune radius decreases toward the Moon, the stability index falls toward unity and the divergence rate slows considerably. A number of stability transitions, or bifurcations, bound the NRHO region: for the L2 family in the Earth–Moon system, these occur at perilune radii of approximately 1 832 km and 17 390 km; for the L1 family, at 934 km and 19 062 km as pointed by the arrows in Figure 2.10 [28, 13].

Within these bounds, both stability indices remain close to unity. To quantify the divergence timescale, a linear time constant τ is defined as the number of orbital revolutions required to amplify a given perturbation by a factor of e [28]:

$$\tau [\text{rev}] = \frac{1}{|\text{Re}(\ln \lambda_{\max}(\Phi(T, 0)))|} \quad (2.7)$$

where $\lambda_{\max}(\cdot)$ is the dominant eigenvalue of the monodromy matrix $\Phi(T, 0)$, and $\tau \rightarrow \infty$ for a marginally stable orbit. While some NRHOs are technically unstable in a linear

constrains the monodromy matrix to be symplectic, forcing its eigenvalues to appear in reciprocal pairs: for every diverging direction ($|\lambda| > 1$) there exists a converging one ($|1/\lambda| < 1$). [13]

sense, the magnitude of their stability index remains close to unity throughout the family [28, 13], implying a slow divergence rate that renders them operationally stable. Larger halo orbits near the Lagrange point, by contrast, possess stability indices much greater than one and depart on sub-revolution timescales, making long-duration maintenance impractical.

Operational Implications

The combination of favourable stability, eclipse geometry, and orbital coverage translates directly into four operationally critical advantages for a long-duration habitat such as the Lunar Gateway.

- **Station-keeping cost.** The slow divergence quantified by τ means the orbit does not require continuous correction. Station-keeping analyses conducted in higher-fidelity ephemeris models confirm that the required ΔV is of the order of a few metres per second per year, achieved with infrequent impulsive manoeuvres near apolune [3]. This is substantially lower than the maintenance cost of classical L2 halo orbits farther from the Moon, whose stability indices are orders of magnitude larger and whose manifold structure drives rapid escape [13].
- **Earth communications.** The elongated, nearly polar geometry of the NRHO ensures that the spacecraft spends the large majority of each revolution at high apolune, continuously visible from Earth-based ground stations [13, 11]. Only during the brief, rapid perilune passage does the spacecraft pass close to the Moon, and even then the line of sight to Earth is not permanently interrupted [28]. Furthermore, the position of L2 beyond the Moon provides nearly continuous line of sight coverage of the lunar far side, enabling relay communications for surface operations that are inaccessible from Earth [13].
- **Eclipse avoidance.** Due to the near-rectilinear shape of the orbit, eclipse events are limited to approximately 60 min per 6.5-day revolution. Moreover, as stated before they can be entirely avoided by selecting the orbital insertion epoch such that the lunar shadow cone aligns with the gaps between successive orbital lobes [13]. This is a significant advantage over lower lunar orbits, where eclipse durations and frequencies impose stringent constraints on power system sizing.
- **Lunar surface access.** Each perilune pass occurs near the lunar poles, providing short but predictable access windows to the lunar south pole region, the primary candidate for a permanent surface outpost [14]. The regularity of the 9:2 synodic resonance means these windows repeat with a fixed pattern every two lunar synodic periods (≈ 59 days), simplifying crew and cargo transfer planning [13].

2.2.5 Lunar Surface and South Pole Access and Constraints

The target PDI for Scenario 2 is located above the lunar South Pole region, consistent with the objective chosen in the Artemis mission planning context [28, 11]. The scientific and strategic interest of this region derives primarily from the confirmed presence of water ice in permanently shadowed craters, which represents a potential resource for in-situ

resource utilisation (ISRU) [29, 30].

The first constraint imposed in the trajectory optimisation is a terminal position constraint: the selenographic latitude of the landing point must satisfy $\phi \leq -80^\circ$. From a geometrical standpoint, this restricts the final state of the descent arc to lie within a spherical cap of angular radius 10° centred on the South Pole. The elongated, near-polar geometry of the southern L2 NRHO makes this region naturally accessible during each perilune pass, without requiring intermediate orbits [28, 11, 30].

The second constraint is a time-of-flight (TOF) upper bound. For crewed lunar surface access from the Gateway, NASA analyses adopt a baseline of 0.5 days (≈ 12 h) each way for polar transfers from the 9:2-resonant NRHO, with an associated ΔV of approximately 700–730 m/s one-way [14, 11]. In this thesis, the TOF constraint is set to a maximum of 1 day, consistent with the 1-day limit adopted in recent direct-transfer studies from the Gateway orbit [30], and slightly more conservative than the $0.5 + 0.05$ day reference value cited in NASA overview documents [14]. This choice avoids convergence to multi-revolution solutions while preserving geometries comparable to those of the Artemis reference mission.

The transfer from the NRHO to the PDI gate is modelled as a three-impulse arc: a departure manoeuvre from the NRHO, a mid-course correction (if required by the optimiser), and a final manoeuvre before the PDI (Powered Descent Initiation). Consistent with the latitude and TOF constraints defined above, the terminal state of the transfer is defined at an altitude of approximately 15 km above the lunar surface, within the South Pole region ($\phi \leq -80^\circ$). The powered descent and landing phases are not modelled in this work, as they require a dedicated treatment - including throttle profiles, terrain-relative navigation, and hazard avoidance - that falls outside the scope of this thesis. The choice of 15 km as the terminal altitude is consistent with the initial conditions adopted in dedicated powered-descent studies [31], to which the present work is intended to be complementary.

3 | Trajectory Optimisation: Problem Formulation




Figure 3.1. Scheme of trajectory optimisation methods. Highlighted nodes (★) indicate the approaches adopted in this work.

3.1 Optimal Control Problem Statement

The trajectory design problem addressed in this thesis can be formulated as an optimal control problem (OCP) where the objective is to find a control history $\mathbf{u}(t)$ and a terminal time t_f that minimise a scalar cost functional J subject to the dynamical constraints of Eq. (2.2) and a set of boundary constraints. In the general continuous-thrust case, $\mathbf{u}(t)$ represents a time-varying thrust vector; in the impulsive approximation adopted in this work, thrust is modelled as a discrete sequence of instantaneous velocity increments $\{\Delta \mathbf{v}_i\}_{i=1}^N$ applied at times $\{t_i\}_{i=1}^N$. This naturally reduces the OCP to a finite-dimensional nonlinear programme (NLP) [32]:

$$\min_{\mathbf{x}} f(\mathbf{x}) \quad \text{subject to} \quad \mathbf{g}(\mathbf{x}) = \mathbf{0}, \quad \mathbf{h}(\mathbf{x}) \leq \mathbf{0} \quad (3.1)$$

where the decision vector $\mathbf{x} = \{t_i, \Delta \mathbf{v}_i\}_{i=1}^N$ ¹ collects the manoeuvre times and velocity increments, $\mathbf{g}(\mathbf{x}) = \mathbf{0}$ enforces the equality constraints (e.g. target state matching at the terminal arc), and $\mathbf{h}(\mathbf{x}) \leq \mathbf{0}$ represents the inequality constraints (e.g. time-of-flight bound, latitude requirement, and transit time through Van Allen belts). The problem is nonlinear because the terminal state depends on the decision variables $\mathbf{x}$ through numerical integration of the equations of motion: a small change in an impulse time or magnitude can produce large, nonlinear changes in the arrival conditions. Both the objective function and the constraint functions are therefore nonlinear - and in general non-convex - in the decision variables. Non-convexity implies the possible existence of multiple local minima, which motivates the use of a global optimisation strategy, as discussed in Section 3.3.1.

The objective function considered throughout this work is the total impulsive ΔV cost:

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

which is directly proportional to the propellant mass fraction through the Tsiolkovsky rocket equation. Minimising J is therefore equivalent to minimising propellant consumption for a spacecraft of fixed specific impulse. For crewed lunar missions, reducing the ΔV budget is critical: it directly increases the mass available for crew systems, life support, and contingency reserves, while reducing launch costs and mission risk [25].

Figure 3.1, adapted from [33], summarises the landscape of trajectory optimisation methods; in the following section each branch is discussed in detail and motivated by the choices made in this work.

3.2 Optimisation Approaches

3.2.1 Thrust Model: Impulsive vs. Finite-Burn

In the impulsive model, velocity increments are applied instantaneously, which is valid when the engine burn duration is short compared to the orbital period. For the transfer arcs considered in this work chemical propulsion burns typically last on the order of minutes, making the impulsive approximation well-justified. Therefore, the model provides a ΔV budget that constitutes a lower bound for the finite-burn case; corrections due to gravity losses during finite burns are typically of the order of 1–5% for chemical systems.

All transfers in this thesis are modelled as impulsive. The finite-burn case, relevant when low-thrust electric propulsion is considered, is left as a natural extension of this work.

¹Further in the work, the decision vector is extended to include the NRHO insertion/departure point as an additional optimisation variable.

3.2.2 Direct vs. Indirect Methods

Indirect methods derive the necessary conditions for optimality from the calculus of variations and the Pontryagin Maximum Principle (PMP). The former provides the mathematical framework for finding functions that extremise integral functionals (such as the total ΔV); the latter extends these conditions to problems with control constraints. Together, they transform the OCP into a two-point boundary value problem (TPBVP): a system of differential equations whose boundary conditions are split between the initial time t_0 and the final time t_f , with a set of unknown initial adjoint variables $\boldsymbol{\lambda}(t_0)$ that must be determined as part of the solution [34, 35]. In the impulsive setting, each intermediate manoeuvre introduces additional interior-point constraints - conditions that must be satisfied at intermediate times $t_i \in (t_0, t_f)$ rather than at the endpoints - further increasing the complexity of the problem [34]. These methods are theoretically rigorous and, when they converge, produce highly accurate solutions. The principal difficulty, however, lies in the adjoint variables $\boldsymbol{\lambda}$ (also known as costate variables): Lagrange multipliers equal in number to the system state variables, which are introduced by the analytical necessary conditions and immediately double the size of the problem [34, 35]. Unlike the state variables, which have direct physical interpretations (position, velocity, mass), the adjoint variables do not admit such an immediate interpretation, although they are not devoid of physical meaning: each adjoint variable is, by construction, the gradient of the performance index with respect to the corresponding state variable. In particular, the adjoint variables associated with velocity form the so-called primer vector², providing some physical intuition on the optimal orientation of the manoeuvres along the trajectory [34, 35]. In practice, however, this intuition offers limited help in constructing an initial guess: the adjoint variables may differ by several orders of magnitude from their corresponding states, and may exhibit discontinuities at the junctions of constrained and unconstrained arcs [34, 35]. Consequently, the TPBVP is notoriously sensitive to initialisation, and this sensitivity is further compounded in the multi-body dynamical environment considered here, where the solution landscape is highly non-convex and accurate initial guesses for the adjoint variables are difficult to obtain [34, 35]. For these reasons, indirect methods are not adopted in this work.

Direct methods avoid the costate formulation entirely by transcribing the OCP into a finite-dimensional nonlinear programme (NLP). For continuous-thrust problems, this is achieved through discretisation of the state and control histories [32, 34, 35]; for impulsive transfers, the control variables are already naturally discrete. The common principle is to convert the optimisation into a parameter search problem, eliminating the shooting approach in favour of simultaneous adjustment of all free parameters [34]. These variables are physical quantities, such as the impulse magnitudes and directions and times of flight, making the problem substantially easier to initialise and to constrain. Since no costate variables appear, the problem size is halved relative to the indirect formulation, and the problematic estimation of initial costates is avoided entirely [34, 35]. Direct transcription³

²The primer vector $\mathbf{p} = \boldsymbol{\lambda}_v$ is defined as the vector of adjoint variables associated with velocity. Its direction indicates the marginal change in the performance index for an instantaneous variation of the velocity at fixed time, and its magnitude and time derivative govern classical optimality conditions for impulsive manoeuvres, such as the requirement that $|\mathbf{p}| = 1$ at an optimal interior impulse and $|\mathbf{p}| \leq 1$ elsewhere along the trajectory [36].

³For continuous-thrust trajectories, *transcription* refers specifically to the discretisation of state and control histories into a finite set of variables at collocation nodes. For impulsive transfers, it denotes the reformulation in terms of these discrete parameters.

schemes also exhibit significantly better robustness, that is, the ability to converge to an optimal trajectory from poor initial guesses [34, 35].

Several families of such schemes exist, such as collocation, pseudospectral, and Runge–Kutta parallel-shooting, each suited to different propulsion regimes [34, 35]. In the cislunar environment, where dynamics are highly sensitive to initial conditions and operational constraints are numerous, these methods are the standard choice for practical mission analysis [13, 5]. The specific scheme adopted in this work is described in the next section.

3.2.3 Trajectory Parametrisation: the Multi-Impulse Approach

Within the class of direct methods introduced in the previous section, the choice of transcription scheme determines how the trajectory is discretised and how continuity is enforced. In the continuous-thrust case this involves replacing control histories with polynomial approximations defined over a discrete time grid [34]; in the impulsive case adopted here, the transcription reduces to selecting a finite set of manoeuvre parameters $\{\Delta \mathbf{v}_i, t_i\}$ that fully characterise the trajectory [32]. Figure 3.2, adapted from [32], illustrates the two main shooting approaches and the parametrisation adopted in this work.




Figure 3.2. *Top:* single shooting propagates the entire arc from one initial condition; numerical errors accumulate over long arcs. *Middle:* multiple shooting introduces free interior nodes $\mathbf{n}_i$ at which defect conditions $\delta_i = \mathbf{x}_{i+1}^- - \mathbf{x}_i^+ = \mathbf{0}$ are enforced as NLP equality constraints, limiting error propagation to individual segments. *Bottom:* the multi-impulse parametrisation adopted in this work; each segment is propagated directly from the terminal state of the preceding one, so continuity is satisfied by construction rather than as an explicit constraint.

In single shooting, the entire trajectory is propagated from a single initial condition; the optimisation variables act only at the departure point. This approach is simple to implement but accumulates numerical errors over long arcs, making it highly sensitive to the initial guess - a critical limitation in the cislunar regime where small perturbations

can lead to qualitatively different trajectories.

In multiple shooting, the trajectory is divided into N segments separated by $N - 1$ free interior nodes $\mathbf{n}_i$. Each node is an independent state variable; propagation is performed separately on each segment, and continuity of the full state is enforced at every junction as an NLP equality constraint (defect condition):

$$\boldsymbol{\delta}_i = \mathbf{x}_{i+1}^- - \mathbf{x}_i^+ = \mathbf{0}, \quad i = 1, \dots, N - 1 \quad (3.3)$$

where $\mathbf{x} = [\mathbf{r}, \mathbf{v}]^\top$ and superscripts $-/+$ denote states immediately before and after node i . If an impulsive manoeuvre $\Delta\mathbf{v}_i$ is applied at node i , the velocity defect is replaced by the ballistic relation

$$\mathbf{v}_i^+ = \mathbf{v}_i^- + \Delta\mathbf{v}_i, \quad (3.4)$$

while the position defect $\mathbf{r}_{i+1}^- - \mathbf{r}_i^+ = \mathbf{0}$ is still enforced. Allowing the solver to temporarily violate the defect during the search — and then drive it to zero at convergence — significantly improves robustness on long, sensitive arcs [32].

In this work, continuity is enforced implicitly through the impulsive $\Delta\mathbf{v}$ parametrisation: each segment is propagated directly from the terminal state of the preceding one, so the defect conditions (3.3) are satisfied by construction and never appear as explicit NLP constraints. The multi-impulse structure retains the key advantage of multiple shooting — limiting propagation to short, individually controlled segments — while delegating the enforcement of continuity to the trajectory propagator rather than to the solver. The specific decision vector structures for Scenario 1 ($N = 4$ impulses, LEO $\rightarrow$ NRHO) and Scenario 2 ($N = 3$ impulses, NRHO $\rightarrow$ PDI) are defined in Chapter 4.

3.2.4 Forward and Backward Propagation

A further degree of freedom in the parametrisation of the multi-impulse scheme is the direction of time integration. Trajectory propagation can be performed forward from the departure state or backward from the arrival state; the choice has practical consequences for the numerical conditioning of the problem and the ease of enforcing boundary conditions.

In forward propagation, the spacecraft state is integrated from an initial condition at t_0 (e.g. a LEO parking orbit) toward the target at t_f (e.g. the NRHO). The boundary condition at arrival must then be enforced as an explicit constraint on the terminal state. In scenario 1, this approach suffers from a fundamental difficulty: the NRHO is a highly unstable orbit, and small deviations in the initial conditions are amplified exponentially along the arc, making it difficult to reach a precise point on the orbit without a very accurate initial guess.

In backward propagation, the integration is initialised at the arrival state — a known point on the NRHO at a specified epoch — and propagated backward in time toward the departure orbit. The arrival boundary condition is therefore satisfied by construction, and the optimiser is only required to satisfy the departure constraint ranges. This reversal exploits the fact that the equations of motion are time-reversible under Newtonian gravity: integrating with a negated time step $-\Delta t$ recovers the physically identical trajectory

traversed in the opposite direction [32].

The practical advantage in the cislunar context is significant. The instability of the NRHO means that trajectories propagated forward from LEO diverge rapidly if the arrival conditions are not met with high precision — a stringent requirement that makes forward shooting ill-conditioned. Backward propagation from a fixed NRHO state decouples the arrival constraint from the search, concentrating the optimisation effort on matching the departure orbit, which is a far less sensitive condition [5].

In scenario 2, conversely, a forward propagation is necessary to exploit the opposite asymmetry: the departure state on the NRHO is precisely known from the reference ephemeris, whereas the PDI gate does not correspond to a fixed target state but to a terminal event.

3.3 Optimisation Algorithms

3.3.1 Particle Swarm Optimisation

Particle Swarm Optimisation (PSO) is a population-based metaheuristic method introduced by Kennedy and Eberhart in 1995 [37]. The algorithm originated from the simulation of social behaviour in animal groups: individuals benefit from sharing information about promising regions of the search space, an advantage analogous to the way birds or fish capitalise on one another’s knowledge when searching for food [37]. In biological flocking models (Figure 3.3, adapted from [38]), individuals respond to three distance-dependent zones: repulsion from nearby neighbours to avoid collision, alignment with intermediate neighbours to maintain coherent motion, and attraction toward distant members to preserve group unity [38]. The PSO algorithm abstracts this complex interaction into a simplified model: particles do not interact directly with each other based on spatial proximity, but instead respond to two global attractors — their personal best position $\mathbf{p}_k$ (cognitive memory) and the swarm’s global best $\mathbf{g}$ (social information sharing) — while maintaining exploratory momentum through an inertia term that plays the role of movement continuity [37].




Figure 3.3. *Left:* Schematic representation of the three distance-dependent zones in biological flocking behaviour: repulsion ($r < r_0$), alignment ($r_0 \leq r < r_1$), and attraction ($r \geq r_1$). *Right:* Corresponding force profiles as a function of inter-individual distance. The lower panel shows a continuous attraction function typical of smooth models.

A swarm of n_p particles explores the d -dimensional decision space $\mathbf{x} \in \mathbb{R}^d$, where d is the number of optimisation variables. At time t each particle k maintains a current position $\mathbf{x}_k$, a velocity $\mathbf{v}_k$, and a memory of its personal best position $\mathbf{p}_k$ (*pbest*). The swarm also tracks the global best position $\mathbf{g}$ (*gbest*) found by any particle. At each iteration, velocities and positions are updated as:

$$\mathbf{v}_k^{(t+1)} = w, \mathbf{v}_k^{(t)} + c_1 r_1! (\mathbf{p}_k - \mathbf{x}_k^{(t)}) + c_2 r_2! (\mathbf{g} - \mathbf{x}_k^{(t)}) \quad (3.5)$$

$$\mathbf{x}_k^{(t+1)} = \mathbf{x}_k^{(t)} + \mathbf{v}_k^{(t+1)} \quad (3.6)$$

where w is the inertia weight, c_1 and c_2 are the cognitive and social acceleration coefficients, and $r_1, r_2 \sim \mathcal{U}(0, 1)$ are independent uniform random scalars drawn at each step [39].

The cognitive term $c_1 r_1 (\mathbf{p}_k - \mathbf{x}_k)$ is called "simple nostalgia" [37], as each particle is attracted toward the best solution it has personally encountered, promoting diversity within the swarm. The social term $c_2 r_2 (\mathbf{g} - \mathbf{x}_k)$ models the sharing of publicised knowledge within the swarm, promoting convergence toward globally promising regions. In the global topology adopted here, $\mathbf{g}$ is the best position ever visited by any particle in the entire swarm; a local topology variant exists in which each particle communicates only with a defined neighbourhood, which can reduce premature convergence at the cost of slower computation [40]. The momentum term $w, \mathbf{v}_k^{(t)}$ is essential: removing it was found to severely impair the algorithm's ability to find global optima, as the stochastic overshooting

induced by the velocity memory allows exploration of unknown regions beyond previously found good solutions [37]. The inertia weight w governs the balance between exploration and exploitation: high values favour exploration, low values favour exploitation of local bests.

In the original formulation, Kennedy and Eberhart set both acceleration constants to 2, so that on average each particle overshoots its target by a factor of 1 per dimension, ensuring thorough coverage of the search space between known good regions [37]. Subsequent work by Shi and Eberhart introduced the inertia weight as a parameter to replace the velocity clamping originally used to prevent divergence, and showed that a linearly decreasing w schedule improves convergence behaviour [39].

PSO requires no gradient information and no initial guess, making it well-suited to the non-convex, multimodal objective landscapes typical of cislunar trajectory design, where multiple local minima correspond to qualitatively different transfer geometries and no reliable first-attempt solution is available a priori [40]. In this work, PSO is implemented via MATLAB’s `particleswarm` function (Global Optimisation Toolbox, R2024b) [41]. The swarm size n_p and the maximum number of iterations are treated as fixed parameters, adjusted via preliminary runs to calibrate the computational cost and the accuracy of the search: a larger swarm provides denser coverage of the search space and reduces the risk of missing the global optimum in highly multimodal landscapes [40]. The *InertiaRange*, *SelfAdjustmentWeight*, and *SocialAdjustmentWeight* are set to their MATLAB defaults, which implement an adaptive inertia weight schedule that automatically balances exploration and exploitation throughout the run [41].

Inequality constraints (for instance minimum perigee altitude or time-of-flight limits) are handled via a penalty function approach. Inequality constraints are less detrimental to algorithm performance than equality constraints: while the latter reduce the degrees of freedom of the problem and drastically narrow the space of admissible solutions, the former merely restrict the feasible region without affecting its dimensionality [40]. Infeasible candidate solutions therefore receive a fictitious high value of the objective function, and their velocity is simultaneously reset to zero so that the subsequent update is driven exclusively by the cognitive and social terms, statistically steering the particle back into the feasible region [40]. This strategy requires no explicit infrastructure that handles constraints in the inner loop and is standard practice for PSO applied to constrained trajectory optimisation [40]. This approach requires no modification to the core PSO update equations: constraint satisfaction is checked entirely within the objective function evaluation, making it straightforward to implement and computationally inexpensive [40]. The global solution obtained from PSO is then used as the initial guess for local refinement.

3.3.2 Sequential Quadratic Programming

While PSO provides global exploration of the search space, it converges slowly, it does not guarantee strict satisfaction of constraints, and might not find the true global optimum. A local gradient-based refinement is therefore necessary to produce the high-accuracy solutions [40]. Sequential Quadratic Programming (SQP) is the standard local refinement method for constrained smooth NLPs in astrodynamics [32, 34]: it iteratively solves

a sequence of quadratic programme (QP) subproblems, each providing a second-order approximation of the original NLP in the neighbourhood of the current iterate [42].

The method is based on the Karush–Kuhn–Tucker (KKT) optimality conditions [33], which generalise the classical unconstrained condition $\nabla\phi = \mathbf{0}$ to problems with constraints. Given a feasible point $\mathbf{x}$ and an assumed active constraint set $\mathbf{c}_a(\mathbf{x}) = \mathbf{0}$, the augmented performance index

$$\tilde{\phi}(\mathbf{x}, \boldsymbol{\lambda}_a) = \phi(\mathbf{x}) + \boldsymbol{\lambda}_a^\top \mathbf{c}_a(\mathbf{x}) \quad (3.7)$$

is introduced, where $\boldsymbol{\lambda}_a$ are the Lagrange multipliers associated with the active constraints. The multipliers act as weights that balance the objective against constraint violations: choosing them appropriately allows the constrained optimisation problem to be changed as the search for a stationary point of $\tilde{\phi}$, which can be treated as an unconstrained problem [33]. Formally, $\boldsymbol{\lambda}_a$ is chosen to nullify the first variation of $\tilde{\phi}$ for any direction $\Delta\mathbf{x}$, not only those tangent to the constraint surface. This leads to the KKT system:

$$\mathbf{g} - \mathbf{G}_a^\top \boldsymbol{\lambda}_a = \mathbf{0}, \quad \mathbf{c}_a(\mathbf{x}) = \mathbf{0}, \quad \lambda_{a,i} \geq 0 \quad \forall i, \quad (3.8)$$

where $\mathbf{g} = \nabla\phi$ is the gradient of the performance index and $\mathbf{G}_a = \partial\mathbf{c}_a/\partial\mathbf{x}$ is the Jacobian of the active constraints. The first equation states that at the optimum the gradient of ϕ must lie in the subspace spanned by the constraint gradients: in geometric terms, $\nabla\phi$ must be perpendicular to the feasible surface, meaning there is no direction along which the objective can be improved without violating at least one constraint [33]. The sign condition $\lambda_{a,i} \geq 0$ ensures that each active constraint is genuinely binding: a positive multiplier indicates that relaxing the i -th constraint would improve ϕ , confirming it is active at the optimum, while a negative value signals that the constraint is unnecessarily included in the active set and should be removed [33].

At each SQP iteration, the correction $\Delta\mathbf{x}$ and updated multipliers $\boldsymbol{\lambda}_a$ are computed by solving the linearised KKT system near the current iterate [33]. The active constraint set is then updated: any inactive constraint violated by the candidate step $\mathbf{x} + \Delta\mathbf{x}$ is added to the active set, and any constraint with $\lambda_{a,i} < 0$ is removed. This active-set strategy ensures that only constraints that are truly active at the solution are enforced, avoiding overconstraining the problem. The iteration continues until the KKT conditions (3.8) are satisfied to within a prescribed tolerance, at which point convergence to a local optimum is guaranteed under standard regularity assumptions [42]. In MATLAB, SQP is available through `fmincon` with the `'sqp'` algorithm option, which handles nonlinear equality and inequality constraints natively and provides robust finite-difference gradient approximation when analytical derivatives are not available [43].

3.3.3 Hybrid PSO-SQP Strategy

Despite their individual merits, neither PSO nor SQP alone is sufficient for the trajectory optimisation problem considered here. The decision space is high-dimensional and non-convex: multiple local minima exist, corresponding to qualitatively different transfer geometries. A purely local method such as SQP converges to the nearest local minimum from the initial guess, with no guarantee of global quality. Conversely, a purely global

method such as PSO provides broad coverage of the solution space but converges slowly near the optimum and typically produces solutions either marginally violate constraints at the final iterate so that could be further refined.

This targeted hybrid PSO $\rightarrow$ SQP strategy is also consistent with the broader literature on hybrid optimisation for astrodynamics [40, 44].

3.4 State of the Art

3.4.1 LEO–NRHO Transfer Literature

The problem of designing transfers from LEO to the NRHO has received growing attention since the selection of the 9:2 southern L_2 NRHO as the baseline orbit for the Lunar Gateway. Early work by Davis et al. [14] characterised the dynamical properties of the Gateway orbit and identified the apolune as the preferred insertion point, confirming that insertion at apolune minimises the required braking ΔV due to the low spacecraft velocity at that orbital phase. Zimovan et al. [13] provided a comprehensive mapping of the NRHO family and their stability indices, while Whitley and Martinez [11] explored the mission design implications for the Artemis architecture.

For crewed transfers specifically, the time-of-flight constraint rules out the use of low-energy families such as Ballistic Lunar Transfers (BLTs) and Weak Stability Boundary (WSB) trajectories, which achieve ΔV savings at the cost of 60–100-day flight times [45]. High-energy direct transfers with flight times of 3–7 days have been studied by multiple groups [13, 14], typically yielding total ΔV budgets in the range of 800–1000 m/s for mid transfer manoeuvres and NRHO insertion burns. A refined total ΔV of approximately 707 m/s was obtained for the direct crewed transfer case using a GA-SNOPT hybrid with multiple shooting in a four-body ephemeris model [5].

3.4.2 NRHO–Lunar Surface Transfer Literature

Whitley and Martinez [11] identified the perilune as the natural departure point for surface access, where the spacecraft is closest to the Moon and orbital velocity is maximum, minimising the ΔV required to initiate the descent. Subsequent work by Zimovan-Spreen et al. [15] characterised the sensitivity of the departure geometry to the NRHO orbital phase and established baseline ΔV estimates for powered descent initiation from the 9:2 resonant southern L_2 orbit.

Recent analyses [30] have explored direct transfers from the NRHO to the lunar South Pole region using a multi-impulse parametrisation in the four-body ephemeris model, demonstrating that descent arcs with time-of-flight constraints under 1 day can be efficiently optimised via hybrid evolutionary-gradient strategies. The terminal phase of the descent - from Powered Descent Initiation (PDI) to touchdown - involves a separate optimisation problem typically solved with direct collocation or polynomial guidance methods [46], which is beyond the scope of this thesis. The present work focuses on the

multi-impulse transfer arc from the NRHO perilune to a low-altitude gate at approximately 15 km above the lunar surface that is consistent with a PDI state, treating the descent initiation as the terminal constraint. The PDI gate altitude is adopted from dedicated powered-descent studies [31], from which the present transfer solutions intend to extract compatible terminal conditions.

3.4.3 PSO in Space Trajectory Optimisation

Particle Swarm Optimisation has been applied to a broad class of space trajectory problems since the early 2000s. Comparative studies [40] have shown that PSO outperforms Genetic Algorithms on several benchmark impulsive transfer problems in terms of convergence speed and solution quality, and is well-suited for use as the global search stage in hybrid frameworks coupled with gradient-based local refinement. A subsequent application in the cislunar regime demonstrated the effectiveness of PSO for hybrid optimisation of low-thrust transfers to the Lunar Gateway [1], confirming that the method’s population-based structure is well-suited to solution spaces characterised by multiple local minima, as typical of multi-body dynamics.

The application of PSO to impulsive LEO–NRHO transfers extends the GA-based framework of [5], substituting the evolutionary algorithm with a more efficient metaheuristic method while retaining the same multi-impulse parametrisation and ephemeris-based dynamics. The second case study—optimisation of the transfer from NRHO to the PDI above the lunar South Pole—addresses a mission phase complementary to [5] and directly connected to the landing scenario of [31].

4 | Methodology and Problem Formulation

This chapter translates the theoretical and mission framework established in Chapters 2 and 3 into concrete computational implementations for the two mission scenarios addressed in this thesis. The two problems share the same dynamical model and optimisation pipeline but differ in their propagation environment, terminal conditions, and constraint sets. Section 4.1 presents the numerical integration setup and the SPICE interface. Sections 4.3–4.4 define the decision vectors, objective functions, and constraint implementations for Scenario 1 (LEO $\rightarrow$ NRHO) and Scenario 2 (NRHO $\rightarrow$ PDI) respectively. Section 4.5 describes the PSO and SQP configurations common to both scenarios, and Section 4.6 summarises the end-to-end optimisation pipeline.

4.1 Dynamical Model Implementation

The four-body equations of motion of Eq. (2.2) are integrated in the J2000 Earth-Centred Inertial (ECI) frame using MATLAB’s `ode113` solver (variable-step and variable-order Adams–Bashforth–Moulton predictor-corrector, orders 1–13), with relative and absolute tolerances of 10^{-9} . These tolerances ensure that position errors remain below $\mathcal{O}(1)$ m over the maximum arc durations of the two scenarios (~ 5 days for Scenario 1, ~ 0.6 days for Scenario 2), as verified by energy conservation checks on ballistic arcs. For Scenario 2 the integrator is additionally configured with `InitialStep` = 60 s and `MaxStep` = 600 s, which guarantee accurate triggering of the PDI altitude gate event introduced in Section 4.4.1. The Earth gravitational parameter is $\mu_{\oplus} = 398\,600.4418 \text{ km}^3/\text{s}^2$, while $\mu_{\mathcal{M}}$ and $\mu_{\odot}$ are provided by the `de440.bsp` kernel [23]. The positions and velocities of the Moon and Sun are retrieved at each integration step via `cspice_spkezr` in the J2000 frame, consistent with the SPICE interface described in Section 2.1.3.

Direction of propagation. The integration direction differs between the two scenarios in accordance with the departure and arrival of each transfer. Scenario 1 trajectories are propagated backward in time from a fixed apolune state on the NRHO toward the departure LEO orbit, as established in Section 3.2.4: the apolune is the natural arrival configuration of the transfer, whereas the LEO state is the outcome of the optimisation and is not known a priori. Scenario 2 trajectories, conversely, are propagated forward in time from a fixed perilune state on the NRHO toward the PDI gate above the lunar south

pole, as established in Section 4.4.1: the perilune is the natural departure configuration, and the arrival epoch is determined by the terminal-event detection at $h = 15$ km.

For Scenario 1, the equivalence between backward and forward physical trajectories follows from the time-reversal symmetry of the four-body equations of motion (2.2) under the substitution $t \rightarrow -t$, $\mathbf{v} \rightarrow -\mathbf{v}$, since no dissipative forces (atmospheric drag, solar radiation pressure) are modelled. The forward mission ΔV budget is therefore recovered exactly from the backward solution by reversing the sign of each velocity increment. For Scenario 2 the propagation is already in mission direction, so no time reversal is required and the impulse signs are retained as integrated.

Impulsive manoeuvres. Impulsive manoeuvres are applied identically in both scenarios as instantaneous velocity increments at segment boundaries, with the position unchanged:

$$\mathbf{v}_i^+ = \mathbf{v}_i^- + \Delta \mathbf{v}_i, \quad \mathbf{r}_i^+ = \mathbf{r}_i^-. \quad (4.1)$$

After each application the integrator is restarted from the updated state at the same epoch. Trajectory continuity is satisfied by construction through sequential propagation, consistent with the multi-impulse transcription of Section 3.2.3.

NRHO fixed-state identification. Each scenario anchors its propagation at a different reference state on the NRHO, identified computationally at run-time from the SPICE kernel `receding_horiz_3189_1burnApo_DiffCorr_15yr.bsp`. For Scenario 1, the apolune — the point of maximum spacecraft–Moon distance along the orbit — is located by scanning $N_{\text{scan}} = 624$ uniformly-spaced states (one every ~ 15 minutes) over one NRHO period $T_{\text{NRHO}} \approx 6.56$ days; the maximum-distance epoch is retained as the apolune state.

For Scenario 2, the same scanning procedure is applied with the minimum-distance state selected as the perilune. In addition, the free- α parametrisation of Case 5 (Section 4.4.4) requires the identification of the two apolunes that bracket the reference perilune, in order to define the apolune-anchored segment `[et_apo_prev, et_apo_next]` over which the phase variable $\alpha \in [0, 1)$ varies. These two apolunes are located by an extended scan of $N_{\text{scan,apo}} = 1248$ samples (one every ~ 30 minutes) over the window $[-1.5 T_{\text{NRHO}}, +1.5 T_{\text{NRHO}}]$ centred on the reference epoch: local maxima of the selenocentric distance are detected, and the closest pair bracketing the perilune is retained. The segment duration $T_{\text{segment}} = \text{et_apo_next} - \text{et_apo_prev}$ and the fractional position of the perilune $\alpha_{\text{peri}} = (\text{et_peri} - \text{et_apo_prev}) / T_{\text{segment}}$ are precomputed once per case and held fixed throughout the optimisation. For the reference epoch of 10 March 2028 the resulting values are $T_{\text{segment}} \approx 6.39$ d and $\alpha_{\text{peri}} \approx 0.4912$, the latter confirming the near-symmetric geometry of the NRHO arc.

4.2 Reference Frames

Five reference frames are employed in this work. The J2000 Earth-Centred Inertial (ECI) frame serves as the universal propagation frame for both scenarios. A lunar orbital plane-aligned frame is used to parametrise the decision variables of Scenario 1, ensuring that fixed numerical bounds on the manoeuvre components correspond to physically consistent directions across different departure epochs. A selenocentric inertial frame is used to

express the impulse manoeuvre components of scenario 2 and to render its trajectory plots. The Moon-fixed IAU_MOON frame is used exclusively to evaluate the latitude constraint at the PDI gate in Scenario 2. A fifth frame, the local RTN frame, is introduced for post-processing and diagnostic analysis of the optimised manoeuvres.

4.2.1 ECI J2000 Frame

All trajectory propagation and constraint evaluation are performed in the J2000 Earth-Centred Inertial (ECI) frame. Its origin is at the Earth’s centre of mass; $\hat{\mathbf{I}}$ points toward the mean vernal equinox at J2000.0, $\hat{\mathbf{K}}$ is aligned with the Earth’s mean rotation axis at the same epoch, and $\hat{\mathbf{J}} = \hat{\mathbf{K}} \times \hat{\mathbf{I}}$ completes the right-handed triad, as illustrated in Figure 4.1. All primary-body ephemerides are provided in this frame via `cspice_spkezr`; the spacecraft position $\mathbf{r}_{sc}$ and velocity $\mathbf{v}_{sc}$ are expressed in J2000 throughout the numerical integration. The J2000 frame as implemented in SPICE is aligned with the International Celestial Reference Frame (ICRF) to within approximately 20 mas¹, a discrepancy negligible for cislunar trajectory design [23].




Figure 4.1. ECI J2000 frame: origin at Earth’s CoM, $\hat{\mathbf{I}}$ toward vernal equinox, $\hat{\mathbf{K}}$ toward north pole, $\hat{\mathbf{J}}$ completing the triad.

4.2.2 Lunar Orbital Plane–Aligned Frame

The decision-variable $\Delta\mathbf{v}$ components are expressed in a frame aligned with the instantaneous lunar orbital plane, as illustrated in Figure 4.2. Because the Earth–Moon line and the orientation of the lunar orbit plane change slowly with epoch, fixed numerical bounds on Cartesian $\Delta\mathbf{v}$ components in J2000 would correspond to slightly different physical directions at different departure dates. By contrast, using a plane-aligned frame keeps the bounds approximately aligned with physically meaningful directions, yielding a more homogeneous PSO search space across different epochs.

¹Milliarcsecond.

The frame is computed at the departure epoch t_{dep} identified by the apolune scan of Eq. (4.6):

$$\hat{\mathbf{X}} = \frac{\mathbf{r}_{\mathcal{M}}}{|\mathbf{r}_{\mathcal{M}}|}, \quad \hat{\mathbf{Z}} = \frac{\mathbf{h}_{\mathcal{M}}}{|\mathbf{h}_{\mathcal{M}}|}, \quad \hat{\mathbf{Y}} = \hat{\mathbf{Z}} \times \hat{\mathbf{X}}, \quad (4.2)$$

where $\mathbf{r}_{\mathcal{M}}$ is the Moon's J2000 position and $\mathbf{h}_{\mathcal{M}} = \mathbf{r}_{\mathcal{M}} \times \dot{\mathbf{r}}_{\mathcal{M}}$ its orbital angular momentum. The rotation matrix $\mathbf{R}$ is computed once at initialisation for runs with a fixed insertion point, and recomputed at each fitness evaluation when the NRHO phase index α is a free variable. The conversion to J2000 before each manoeuvre application is:

$$\Delta \mathbf{v}_i^{\text{J2000}} = \mathbf{R}^\top \Delta \mathbf{v}_i^{\text{frame}}, \quad i \in \{\text{NRI}, \text{mid1}, \text{mid2}\}. \quad (4.3)$$




Figure 4.2. Lunar orbital plane-aligned frame: $\hat{\mathbf{X}}$ Earth-to-Moon, $\hat{\mathbf{Z}}$ normal to lunar plane, $\hat{\mathbf{Y}}$ prograde.

4.2.3 Selenocentric Inertial Frame

A selenocentric inertial frame, illustrated in Figure 4.3, is introduced to support the geometric interpretation of Scenario 2 results: the components of the impulsive manoeuvres in the decision vector are expressed in this frame, and all three-dimensional trajectory plots are rendered in it. The frame is defined by translating the origin from the Earth to the Moon and aligning the axes with the Earth–Moon geometry at the reference perilune epoch $t_{\text{peri,ref}}$:

$$\hat{\mathbf{X}}_{\text{sel}} = \frac{\mathbf{r}_{\mathcal{M}}(t_{\text{peri,ref}})}{\|\mathbf{r}_{\mathcal{M}}(t_{\text{peri,ref}})\|}, \quad \hat{\mathbf{Z}}_{\text{sel}} = \frac{\mathbf{r}_{\mathcal{M}}(t_{\text{peri,ref}}) \times \mathbf{v}_{\mathcal{M}}(t_{\text{peri,ref}})}{\|\mathbf{r}_{\mathcal{M}}(t_{\text{peri,ref}}) \times \mathbf{v}_{\mathcal{M}}(t_{\text{peri,ref}})\|}, \quad \hat{\mathbf{Y}}_{\text{sel}} = \hat{\mathbf{Z}}_{\text{sel}} \times \hat{\mathbf{X}}_{\text{sel}} \quad (4.4)$$

where $\mathbf{r}_{\mathcal{M}}$ and $\mathbf{v}_{\mathcal{M}}$ are the Moon position and velocity in J2000 at the reference epoch, retrieved from the SPICE kernel `de440.bsp`. The unit vector $\hat{\mathbf{X}}_{\text{sel}}$ thus points outward along the Earth–Moon line (Earth-to-Moon), $\hat{\mathbf{Z}}_{\text{sel}}$ is aligned with the instantaneous lunar orbital angular momentum, and $\hat{\mathbf{Y}}_{\text{sel}}$ completes the right-handed triad, oriented along the prograde direction of the Moon's orbital motion. The rotation matrix $\mathbf{R}_{\text{sel}} = [\hat{\mathbf{X}}_{\text{sel}} \ \hat{\mathbf{Y}}_{\text{sel}} \ \hat{\mathbf{Z}}_{\text{sel}}]^\top$

is computed once at initialisation and held fixed for the entire optimisation run, including the Case 5 free- α configuration in which the departure epoch t_{dep} varies with the decision variable α : the frame is anchored to the reference perilune, not to the moving departure point, so that the interpretation of the impulse components remains consistent across the multi-start campaign.

The frame is inertial (non-rotating) at every epoch within a given optimisation run, despite being aligned with a configuration of the Earth-Moon geometry that evolves over the mission duration. This distinction is significant: the IAU_MOON frame (Section 4.2.4) co-rotates with the lunar body and is used only at the terminal-event epoch to evaluate the selenographic latitude φ and altitude h_{PDI} . The selenocentric inertial frame defined here, by contrast, provides a geometrically intuitive but kinematically inertial reference for expressing impulse directions and visualising the descent geometry. Trajectory states are transformed from J2000 into this frame for post-processing only; all numerical propagation remains in the J2000 inertial frame as discussed in Section 4.1.




Figure 4.3. Selenocentric inertial frame: $\hat{\mathbf{X}}_{\text{sel}}$ along the Earth-to-Moon outward direction (pointing away from the Earth as seen from the Moon), $\hat{\mathbf{Z}}_{\text{sel}}$ aligned with the Moon’s instantaneous orbital angular momentum, and $\hat{\mathbf{Y}}_{\text{sel}}$ completing the right-handed triad along the prograde direction of the lunar orbital motion. The Earth is shown schematically (not to scale).

4.2.4 Body-Fixed IAU_MOON Frame

The Moon-fixed IAU_MOON frame, shown in Figure 4.4, is a body-fixed, co-rotating frame used exclusively to evaluate the selenographic latitude and altitude at the PDI event in Scenario 2. Unlike the three frames introduced above, whose rotation matrices are either fixed once per run or updated only at the scale of the multi-start campaign, the IAU_MOON frame rotates continuously with the Moon’s physical rotation and is therefore evaluated at the specific epoch of interest - here, the terminal event epoch t_{PDI} - rather than held constant.

The frame is defined by the IAU orientation model for the Moon, which specifies the right ascension α_0 and declination δ_0 of the lunar rotation pole and the prime-meridian angle W as functions of time, tabulated in the SPICE generic lunar kernel `pck00011.tpc`.

The distinction between $\hat{\mathbf{Z}}_{\text{IAU}}$ and the selenocentric-inertial frame's $\hat{\mathbf{Z}}_{\text{sel}}$ (Section 4.2.3) is that the true lunar rotation pole is offset from the Moon's orbital angular momentum vector by approximately 6.7° . Evaluating the selenographic latitude against $\hat{\mathbf{Z}}_{\text{sel}}$ instead of the true rotation pole would therefore introduce an error of the same order, operationally unacceptable given the tightness of the -85° South Pole constraint of Case 4. The IAU_MOON frame is used only at the terminal PDI epoch for this reason; all other Scenario 2 quantities (impulse components, trajectory visualisation) are expressed in the inertial selenocentric frame.




Figure 4.4. Body-fixed IAU_MOON frame: $\hat{\mathbf{X}}_{\text{IAU}}$ along the mean prime meridian, $\hat{\mathbf{Z}}_{\text{IAU}}$ along the true lunar rotation pole, and $\hat{\mathbf{Y}}_{\text{IAU}}$ completing the right-handed triad. The dashed vector marks the lunar orbital angular momentum direction $\hat{\mathbf{Z}}_{\text{sel}}$, offset from the true rotation pole by approximately 6.7° .

4.2.5 Local RTN Frame

The RTN (Radial–Transverse–Normal) frame is defined at each spacecraft state $(\mathbf{r}, \mathbf{v})$ as:

$$\hat{\mathbf{R}} = \frac{\mathbf{r}}{\|\mathbf{r}\|}, \quad \hat{\mathbf{N}} = \frac{\mathbf{r} \times \mathbf{v}}{\|\mathbf{r} \times \mathbf{v}\|}, \quad \hat{\mathbf{T}} = \hat{\mathbf{N}} \times \hat{\mathbf{R}}. \quad (4.5)$$

$\hat{\mathbf{R}}$ points radially outward, $\hat{\mathbf{N}}$ is aligned with the instantaneous orbital angular momentum, and $\hat{\mathbf{T}}$ is tangential in the direction of motion, as shown in Figure 4.5. Decomposing an impulsive manoeuvre $\Delta \mathbf{v}$ into RTN components provides a direct physical interpretation: Δv_R modifies eccentricity and argument of perigee, Δv_T changes the orbital energy, and Δv_N drives out-of-plane manoeuvres. This frame is used exclusively for post-processing and it does not enter the optimisation formulation.




Figure 4.5. Local RTN frame at the spacecraft state $(\mathbf{r}, \mathbf{v})$: $\hat{\mathbf{R}}$ radial, $\hat{\mathbf{T}}$ tangential, $\hat{\mathbf{N}}$ normal. Used for post-processing of optimised manoeuvres only.

4.3 Scenario 1: LEO $\rightarrow$ NRHO

4.3.1 Mission Definition

Scenario 1 models a crewed transfer from a circular LEO parking orbit to the 9:2 southern L_2 NRHO. The nominal arrival epoch at the NRHO is set to 2028-03-10T00:00:00 UTC, a representative date within the Artemis IV operational window [47]. The actual departure epoch t_{dep} is determined by scanning 624 uniformly spaced states over one NRHO period ($T_{\text{NRHO}} = 6.56$ days) starting from the arrival epoch, and selecting the state of maximum spacecraft–Moon distance:

$$t_{\text{dep}} = \arg \max_{t \in [t_{\text{arr}}, t_{\text{arr}} + T_{\text{NRHO}}]} \|\mathbf{r}_{\text{sc}}(t) - \mathbf{r}_{\mathcal{M}}(t)\|, \quad (4.6)$$

where $\mathbf{r}_{\text{sc}}(t)$ and $\mathbf{r}_{\mathcal{M}}(t)$ are the spacecraft and Moon position vectors in J2000, both retrieved via `cspice_spekezt` at each scan epoch. For the nominal epoch of 10 March 2028, the scan identifies the apolune at $t_{\text{dep}} = t_{\text{arr}} + 0.8539 T_{\text{NRHO}}$, corresponding to a spacecraft–Moon distance of approximately 71 082 km. The backward propagation is initialised at t_{dep} with the NRHO state $\mathbf{s}_{\text{NRHO}}(t_{\text{dep}})$ retrieved from the `receding_horiz_3189_1burnApo_DiffCorr_15yr.bsp` kernel. The departure orbit constraints - altitude range, inclination window, and Van Allen belt transit limit - are translated into numerical penalty terms in Section 4.3.3.

4.3.2 Decision Vector

A sequence of $N = 3$ impulsive manoeuvres is admitted, yielding a 12-dimensional decision vector:

$$\mathbf{x}_1 = \left[\underbrace{\Delta v_{\text{NRI},x}, \Delta v_{\text{NRI},y}, \Delta v_{\text{NRI},z}}_{\Delta \mathbf{v}_{\text{NRI}}}, \tau_1, \underbrace{\Delta v_{\text{mid1},x}, \Delta v_{\text{mid1},y}, \Delta v_{\text{mid1},z}}_{\Delta \mathbf{v}_{\text{mid1}}}, \tau_2, \underbrace{\Delta v_{\text{mid2},x}, \Delta v_{\text{mid2},y}, \Delta v_{\text{mid2},z}}_{\Delta \mathbf{v}_{\text{mid2}}}, \tau_3 \right]^\top \in \mathbb{R}^{12} \quad (4.7)$$

where $\Delta \mathbf{v}_{\text{NRI}}$ is the NRHO departure manoeuvre applied at the apolune, $\Delta \mathbf{v}_{\text{mid1}}$ and $\Delta \mathbf{v}_{\text{mid2}}$ are mid-course correction manoeuvres, and τ_1, τ_2, τ_3 are the corresponding segment durations in days. The PSO optimises the $\Delta \mathbf{v}$ components as expressed in the lunar orbital plane frame introduced in Section 4.2; before each fitness evaluation, these components are rotated into the J2000 ECI frame via Eq. (4.3) for trajectory propagation.

The bounds on the decision vector are summarised in Table 4.1. The bounds on $\Delta \mathbf{v}_{\text{NRI}}$ are asymmetric: the x -component (approximately radial with respect to the Moon) is restricted to $[-1.20, +0.30]$ km/s, reflecting the predominantly retrograde character of the departure burn in the lunar frame. The mid-course components are bounded symmetrically at ± 0.30 km/s, consistent with the small correction magnitudes expected for this transfer class. The segment duration bounds were progressively relaxed during the tuning campaign: an initial configuration with $\tau_{\min} = 0.30$ days and $\tau_{\max} = 3$ days caused the PSO solutions to consistently saturate the lower bounds, indicating an artificially constrained search space. Reducing the minimum to $\tau_{\min} = 0.05$ days partially alleviated the issue; setting $\tau_{\min} = 0$ and $\tau_{\max} = 10$ days in the final configuration allowed the optimiser to implicitly collapse degenerate segments to zero duration while accommodating the full range of geometrically feasible transfer arcs. The operational time-of-flight limit of 5 days is subsequently enforced through the penalty function of Section 4.3.3, rather than through the variable bounds.

Table 4.1. Decision variable bounds for Scenario 1 (LEO $\rightarrow$ NRHO). All velocity components are expressed in the lunar orbital plane-aligned frame of Section 4.2.

Variable	Lower bound	Upper bound	Units
$\Delta v_{\text{NRI},x}$	-1.20	+0.30	km/s
$\Delta v_{\text{NRI},y}$	-0.80	+0.80	km/s
$\Delta v_{\text{NRI},z}$	-0.30	+0.30	km/s
τ_1	0	10	days
$\Delta v_{\text{mid1},\{x,y,z\}}$	-0.30	+0.30	km/s
τ_2	0	10	days
$\Delta v_{\text{mid2},\{x,y,z\}}$	-0.30	+0.30	km/s
τ_3	0	10	days

4.3.3 Objective Function and Constraint Implementation

The objective functional for Scenario 1 is the total impulsive ΔV cost, including the three transfer manoeuvres and the LEO insertion burn:

$$J_1(\mathbf{x}_1) = \|\Delta \mathbf{v}_{\text{NRI}}\| + \|\Delta \mathbf{v}_{\text{mid1}}\| + \|\Delta \mathbf{v}_{\text{mid2}}\| + \Delta v_{\text{circ}}, \quad (4.8)$$

where Δv_{circ} denotes the LEO insertion burn evaluated at the perigee of the transfer arc, computed as the scalar difference between the spacecraft speed and the local circular velocity:

$$\Delta v_{\text{circ}} = \left| v_{\text{peri}} - \sqrt{\frac{\mu_{\oplus}}{r_{\text{peri}}}} \right|. \quad (4.9)$$

This expression is valid when the velocity at perigee is predominantly tangential, which follows from the perigee detection condition $\mathbf{r} \cdot \mathbf{v} = 0$, which is a terminal condition for all solutions.

Inequality constraints are enforced through a penalty function approach, in which each violated constraint adds a term proportional to its residual to the objective:

$$\tilde{J}_1(\mathbf{x}_1) = J_1(\mathbf{x}_1) + \sum_k \rho_k \max(0, c_k(\mathbf{x}_1)), \quad (4.10)$$

where c_k is the k -th constraint residual and ρ_k is the corresponding penalty weight. The penalty weights were determined empirically through preliminary tuning runs, with the guiding criterion that a typical constraint violation produces a penalty of order $\mathcal{O}(1)$ km/s, comparable to the nominal transfer ΔV . The active constraint set varies across the five cases and each case activates all constraints of the preceding one plus one additional term, as summarised in Table 4.2.

Table 4.2. Constraint implementation for Scenario 1.

Constraint	Residual $c_k \leq 0$	Weight ρ_k
Perigee altitude (lower)	$150 - h_{\text{peri}}$	2.0
Perigee altitude (upper)	$h_{\text{peri}} - 600$	2.0
Total time of flight	$\tau_{\text{tot}} - 5$	0.5
Inclination, Cape Canaveral (lower)	$26 - i$	3.0
Inclination, Cape Canaveral (upper)	$i - 31$	3.0
Inclination, Kourou (lower)	$4 - i$	3.0
Inclination, Kourou (upper)	$i - 8$	3.0
Van Allen inner belt transit	$t_{\text{inner}} - 0.5$	1.5
Van Allen outer belt transit	$t_{\text{outer}} - 2.0$	1.5
Earth impact	$-h_{\text{peri}}$ if $h_{\text{peri}} < 0$	10.0

The inclination constraints for Cape Canaveral and Kourou are mutually exclusive and are activated individually for Cases 3a and 3b respectively. The Van Allen constraint is activated for Case 4 (with operational and tightened-threshold sub-variants, Cases 4a-4d) and retained for the subsequent free-arrival-point case (Case 5).

The constraint set described above assumes that the backward propagation is initialised at the apolune, i.e. the NRHO phase index α is fixed at α_{apo} throughout Cases 1–5. A further extension of the problem relaxes this assumption by admitting α as a 13th free variable in the decision vector, allowing the optimiser to explore departure points across the entire NRHO period. In this extended formulation the objective and penalty structure of Eq. (4.10) remain unchanged; only the initialisation of the backward propagation and the computation of the rotation matrix $\mathbf{R}$ are affected.

4.4 Scenario 2: NRHO $\rightarrow$ Powered Descent Initiation

4.4.1 Mission Definition

Scenario 2 models a crewed descent transfer from the 9:2 southern L_2 NRHO to a terminal state compatible with the Powered Descent Initiation (PDI) phase of a lunar South Pole landing. The trajectory is parametrised with $N = 3$ impulses: a departure manoeuvre $\Delta\mathbf{v}_{\text{dep}}$ applied at the NRHO perilune, a single mid-course correction $\Delta\mathbf{v}_{\text{mcc}}$, and a final braking manoeuvre $\Delta\mathbf{v}_{\text{PDI}}$ delivered immediately before crossing the PDI gate at an altitude of approximately 15 km above the mean lunar surface. The powered descent and landing phases below this gate involve a qualitatively different guidance problem (throttle scheduling, terrain-relative navigation, hazard avoidance) and are not modelled in this work. The converged solutions are intended to provide compatible upstream initial conditions for dedicated powered-descent studies.

Integration direction. In contrast to Scenario 1, the trajectory is propagated forward in mission time, from the NRHO perilune toward the PDI gate. This is not a deviation from the strategy presented previously but a direct consequence of it: the integration in both scenarios proceeds from the fully-specified boundary (NRHO) toward the partially-specified one (PDI). In Scenario 1 the NRHO arrival state is uniquely defined by the ephemeris kernel, while the LEO departure is constrained only by an altitude window and an inclination band, so the integration runs backward in mission time from NRHO to LEO. In Scenario 2 the NRHO perilune is again the fully-specified boundary, whereas the PDI gate imposes only an altitude condition and a latitude bound (in the constrained cases), leaving the longitude and the precise arrival epoch as free quantities. Treating PDI as an initial condition for backward integration would require parameterising these six unspecified terminal degrees of freedom as additional decision variables, increasing the dimensionality of the search space from 12 to 18 without yielding physical insight. Forward integration from the perilune avoids this overhead: the optimiser shapes the descent arc through the manoeuvre parameters alone, and the terminal state at the PDI gate is recovered implicitly. The two formulations are mathematically equivalent under the time-reversal symmetry, so the converged ΔV budget is unaffected by the choice; the rationale is therefore one of search-space economy alone.

Reference epoch and departure state. The reference NRHO arrival epoch is set to $t_{\text{ref}} = 2028\text{-}03\text{-}10\text{T}00\text{:}00\text{:}00$ UTC, the same nominal date adopted for Scenario 1. The perilune epoch within the corresponding NRHO period is identified by a 3000-point uniform

scan over $[t_{\text{ref}}, t_{\text{ref}} + T_{\text{NRHO}}]$ with $T_{\text{NRHO}} = 6.56$ days, selecting the epoch of minimum spacecraft–Moon distance:

$$t_{\text{dep}} = \arg \min_{t \in [t_{\text{ref}}, t_{\text{ref}} + T_{\text{NRHO}}]} \|\mathbf{r}_{\text{sc}}(t) - \mathbf{r}_M(t)\|. \quad (4.11)$$

The spacecraft and Moon position vectors $\mathbf{r}_{\text{sc}}(t)$ and $\mathbf{r}_M(t)$ are retrieved in J2000 via `cspice_spkezr` from the `receding_horiz_3189_1burnApo_DiffCorr_15yr.bsp` and `de440.bsp` kernels respectively. For the nominal arrival epoch, the scan identifies a perilune altitude $h_{\text{peri}}^{\text{NRHO}} \approx 1636$ km, consistent with the NRHO characterisation in Section 2.2.4. The perilune is the natural departure point for surface access, as established by Whitley and Martinez [11]: at this point the spacecraft is closest to the Moon and the orbital velocity is maximum, minimising the impulse required to initiate the descent.

Dynamical setting. Unlike Scenario 1, the descent arc lies entirely within the lunar gravitational sphere of influence ($r_{\text{SOI},M} \approx 66\,200$ km) for the majority of its duration. The Moon’s gravity is therefore the dominant attractor and the Sun’s contribution is secondary, although both are retained in the four-body model of Eq. (2.2). The integration is performed in the J2000 ECI frame with `ode113` at the tolerances established in Section 4.1. To accelerate the inner loop of the optimisation, the Moon and Sun ephemerides are pre-tabulated on a uniform temporal grid with $\Delta t_{\text{grid}} = 120$ s steps spanning the maximum admissible TOF, and the third-body states required by the right-hand side of Eq. (2.2) are obtained at each integration step by linear interpolation between adjacent grid points. This substitution incurs an interpolation error on the lunar position below 50 m over the entire propagation window, well within the integrator tolerance, and reduces the per-evaluation cost by replacing repeated SPICE queries with $\mathcal{O}(1)$ array look-ups.

Terminal condition. The descent arc is terminated by an event-location procedure that detects the spacecraft crossing the PDI gate altitude on the descending direction. Defining the altitude above the mean lunar surface as $h(t) = \|\mathbf{r}_{\text{sc}}(t) - \mathbf{r}_M(t)\| - R_M$ with $R_M = 1737.4$ km, the terminal event is

$$g(t) = h(t) - h_{\text{PDI}} = 0, \quad h_{\text{PDI}} = 15 \text{ km}, \quad \dot{h}(t) < 0, \quad (4.12)$$

implemented via the event facility of `ode113` (terminal, descending direction) during integration of the third segment. The descending condition $\dot{h} < 0$ filters out spurious crossings during the early phases of the arc, where the trajectory may temporarily approach the 15 km altitude from below before drifting upward. A secondary safety event monitors a hard impact with the lunar surface ($h(t) = 0$), which is treated as a constraint violation rather than a valid terminus.

4.4.2 Decision Vector

The three-impulse parametrisation yields a 12-dimensional decision vector:

$$\mathbf{x}_2 = \left[\underbrace{\Delta v_{\text{dep},x}, \Delta v_{\text{dep},y}, \Delta v_{\text{dep},z}}_{\Delta \mathbf{v}_{\text{dep}}}, \tau_1, \underbrace{\Delta v_{\text{mcc},x}, \Delta v_{\text{mcc},y}, \Delta v_{\text{mcc},z}}_{\Delta \mathbf{v}_{\text{mcc}}}, \tau_2, \underbrace{\Delta v_{\text{PDI},x}, \Delta v_{\text{PDI},y}, \Delta v_{\text{PDI},z}}_{\Delta \mathbf{v}_{\text{PDI}}}, \tau_3 \right]^\top \in \mathbb{R}^{12} \quad (4.13)$$

where $\Delta \mathbf{v}_{\text{dep}}$ is the departure manoeuvre applied at perilune, $\Delta \mathbf{v}_{\text{mcc}}$ is the mid-course correction, $\Delta \mathbf{v}_{\text{PDI}}$ is the final braking manoeuvre, and τ_1, τ_2, τ_3 are the durations of the corresponding ballistic segments, expressed in days. All impulse components are expressed in the selenocentric frame of Eq. (4.4) and are rotated into J2000 via $\mathbf{R}_{\text{sel}}^\top$ before each integration step. The bounds adopted for the decision variables are reported in Table 4.3.

Table 4.3. Decision variable bounds for Scenario 2 (NRHO $\rightarrow$ PDI). All velocity components are expressed in the selenocentric frame of Eq. (4.4).

Variable	Lower bound	Upper bound	Units
$\Delta v_{\text{dep},\{x,y,z\}}$	-1.00	+1.00	km/s
τ_1	0.05	0.60	days
$\Delta v_{\text{mcc},\{x,y,z\}}$	-0.30	+0.30	km/s
τ_2	0.01	0.40	days
$\Delta v_{\text{PDI},\{x,y,z\}}$	-1.00	+1.00	km/s
τ_3	0.01	0.20	days

The bounds on $\Delta \mathbf{v}_{\text{dep}}$ and $\Delta \mathbf{v}_{\text{PDI}}$ are deliberately wide and symmetric, reflecting the absence of a preferred direction for either the NRHO departure impulse or the pre-PDI braking. Preliminary parametric exploration indicates that South Pole access is driven predominantly by an out-of-plane component ($-\hat{\mathbf{z}}_{\text{sel}}$) of the departure burn, supplemented by a retrograde in-plane contribution; the symmetric ± 1.00 km/s envelope accommodates these directions while preserving the optimiser’s freedom to explore alternative geometries. The mid-course correction is bounded symmetrically at ± 0.30 km/s, consistent with the small correction magnitudes expected for this transfer class. The segment-duration bounds reflect the descent geometry: τ_1 covers the bulk of the arc from perilune toward lunar approach, τ_2 accommodates the mid-course correction window, and τ_3 provides a short pre-PDI segment within which the terminal event of Eq. (4.12) is expected to trigger. The total-TOF operational limit ($\tau_{\text{tot}} \leq 0.6$ days) is enforced through the penalty function of the next subsection rather than through the variable bounds, so that the optimiser may explore TOF-relaxed solutions during the exploratory phase before converging within the operational window.

4.4.3 Objective Function and Constraint Implementation

The objective function is the total on-board impulsive cost of the descent phase:

$$J_2(\mathbf{x}_2) = \|\Delta \mathbf{v}_{\text{dep}}\| + \|\Delta \mathbf{v}_{\text{mcc}}\| + \|\Delta \mathbf{v}_{\text{PDI}}\|. \quad (4.14)$$

No surface-insertion term is included, as the trajectory terminates ballistically at the PDI gate and the subsequent powered descent is out of scope.

The constraint set for Scenario 2 differs substantially from that of Scenario 1: arrival at the PDI gate at the prescribed altitude and inertial velocity, latitude within the South Pole region, and compliance with a significantly tighter time-of-flight cap. Inequality constraints are enforced through the same penalty-augmented objective formalism adopted in Section 4.3.3:

$$\tilde{J}_2(\mathbf{x}_2) = J_2(\mathbf{x}_2) + \sum_k \rho_k \max(0, c_k(\mathbf{x}_2)), \quad (4.15)$$

where c_k is the k -th constraint residual and ρ_k is the corresponding penalty weight. The residuals and weights adopted for Scenario 2 are summarised in Table 4.4.

Table 4.4. Constraint implementation for Scenario 2.

Constraint	Residual c_k	Weight ρ_k
PDI gate not reached (lower)	$10 - h_{\text{PDI}}^* \leq 0$	0.5
PDI gate not reached (upper)	$h_{\text{PDI}}^* - 20 \leq 0$	0.5
PDI velocity window (lower)	$1.660 - v_{\text{PDI}} \leq 0$	10.0
PDI velocity window (upper)	$v_{\text{PDI}} - 1.710 \leq 0$	10.0
Total time of flight	$\tau_{\text{tot}} - 0.6 \leq 0$	2.0
Southern hemisphere (Case 3)	$\phi_{\text{PDI}} - 0 \leq 0$	1.0
South Pole region (Case 4)	$\phi_{\text{PDI}} + 85 \leq 0$	1.0
Lunar surface impact	$-h_{\text{min}}$ if $h_{\text{min}} < 0$	10.0

The PDI altitude window $[10, 20]$ km in Table 4.4 acts as a soft fall-back when the event of Eq. (4.12) fails to trigger within the allotted duration of the third segment: in that case the residual is evaluated on the terminal altitude h_{PDI}^* at the end of integration, driving the optimiser toward arcs that at least graze the prescribed band. When the event does trigger, the altitude is exactly $h_{\text{PDI}} = 15$ km by construction and both upper and lower residuals are inactive.

The selenographic latitude φ_{PDI} at the PDI event is computed by transforming the spacecraft position from J2000 into the body-fixed IAU_MOON frame via `cspice_spkezr`, and evaluating the latitude angle from the Moon-centred position vector:

$$\varphi_{\text{PDI}} = \arcsin\left(\frac{z_{\text{bf}}}{\|\mathbf{r}_{\text{sc}}^{\mathcal{M}}\|}\right)\bigg|_{t=t_{\text{PDI}}}, \quad (4.16)$$

where $\mathbf{r}_{\text{sc}}^{\mathcal{M}} = (x_{\text{bf}}, y_{\text{bf}}, z_{\text{bf}})$ is the spacecraft position in IAU_MOON coordinates at the event time. The use of the Moon-fixed frame ensures that the latitude is measured with

respect to the lunar rotational pole, consistent with the conventional definition adopted for South Pole landing-site characterisation.

Cumulative constraint cases. The cases analysed are introduced cumulatively, each adding one new requirement or one additional degree of freedom to the preceding set. The total time of flight $\tau_{\text{tot}} \leq 0.6$ d is introduced as an operational constraint from Case 4 onward, once the south-polar latitude requirement is activated, and is retained in the subsequent cases; it is not enforced in Cases 1-3, in which the descent’s natural time of flight is left unconstrained.

Case 1 - Single-impulse baseline. A single impulsive manoeuvre at the reference perilune places the spacecraft on a ballistic arc to the PDI gate. The altitude window ($h^* = 15$ km), the velocity window ($|v_{\text{PDI}}| \in [1.660, 1.710]$ km/s), the surface-impact safety term and the TOF bound $\tau_{\text{tot}} \leq 0.6$ d are all active. The decision vector has dimension $n_{\text{vars}} = 4$ ($\Delta \mathbf{v}_{\text{dep}} \in \mathbb{R}^3, \tau_1$).

Case 2 — Three-impulse parametrisation. Same constraint set as Case 1, but with the decision vector extended to $n_{\text{vars}} = 12$ to admit three impulses. The two intermediate impulses $\Delta \mathbf{v}_{\text{mcc}}$ and $\Delta \mathbf{v}_{\text{PDI}}$ and the additional segment durations τ_2, τ_3 enlarge the search space without modifying the active-constraint set, so that the marginal ΔV benefit of the three-impulse parametrisation per se can be isolated. The PDI arrival latitude remains unconstrained.

Case 3 — Southern hemisphere. Same as Case 2, with the additional requirement $\varphi_{\text{PDI}} \leq 0^\circ$. The constraint isolates the marginal ΔV cost of biasing the descent geometry toward the southern hemisphere.

Case 4 — South Pole region. Same as Case 3, with the southern-hemisphere bound replaced by the stricter $\varphi_{\text{PDI}} \leq -85^\circ$, and with the operational time-of-flight cap $\tau_{\text{tot}} \leq 0.6$ d introduced jointly, retained in the subsequent cases.

Case 5 — Free departure point on the NRHO. Same constraint set as Case 4, with the NRHO phase α admitted as a thirteenth optimisation variable (Section 4.4.4); $n_{\text{vars}} = 13$. The optimiser is free to relocate the departure point along the apolune-anchored segment $[t_{\text{apo,prev}}, t_{\text{apo,next}}]$ that brackets the reference perilune, rather than departing strictly at the perilune itself.

The progression mirrors the cumulative logic adopted for Scenario 1: each case retains all active constraints of the previous one and introduces either a single new requirement (Cases 3, 4) or a single additional degree of freedom (Cases 2, 5). The mutually exclusive activation of the southern-hemisphere and South-Pole latitude rows in Table 4.4 reflects the cumulative-replacement structure: in Case 4 the South-Pole bound supersedes the southern-hemisphere bound, which is automatically satisfied by any solution meeting $\varphi_{\text{PDI}} \leq -85^\circ$.

Penalty-weight calibration. The penalty weights of Table 4.4 were determined empirically through preliminary tuning runs, following the same guiding criterion adopted for Scenario 1: a typical constraint violation should produce a penalty contribution of order $\mathcal{O}(1)$ km/s, comparable to the nominal descent ΔV budget of ~ 1 km/s. The TOF weight is set higher than in Scenario 1 ($\rho_{\text{TOF}} = 2.0$ versus 0.5) to reflect the considerably tighter operational limit (0.6 days versus 5 days), while the latitude weight ($\rho_{\text{lat}} = 1.0$ per degree

of overshoot) accommodates the wider residual range expected when the optimiser explores high-latitude geometries from a near-equatorial approach. The surface-impact weight ($\rho_{\text{imp}} = 10.0$) is deliberately the stiffest term in the table, mirroring the Earth-impact safeguard of Scenario 1, and ensures that no converged solution intersects the lunar surface during any segment of the arc.

The constraint set described above assumes that the forward propagation is initialised at the NRHO perilune, i.e. the NRHO phase index α is fixed at α_{peri} throughout Cases 1–4. A further extension of the problem, described in the next subsection, relaxes this assumption by admitting α as a 13th free variable in the decision vector, allowing the optimiser to explore departure points across the entire NRHO period.

4.4.4 Extension: Free Departure Point on the NRHO

In Cases 1–4 the backward boundary condition is fixed at the NRHO perilune, exploiting the geometrical argument that the perilune minimises the departure impulse for a given target altitude. This restriction can be removed by promoting the NRHO phase along the orbit to a decision variable, in close analogy with the equivalent extension of Scenario 1. The decision vector is augmented to thirteen components:

$$\mathbf{x}'_2 = [\mathbf{x}_2^\top, \alpha]^\top \in \mathbb{R}^{13}, \quad \alpha \in [0, 1), \quad (4.17)$$

where α is the normalised orbital phase along one NRHO period relative to the reference perilune epoch:

$$t_{\text{dep}}(\alpha) = t_{\text{ref}} + \alpha T_{\text{NRHO}}. \quad (4.18)$$

At each fitness evaluation the NRHO state corresponding to $t_{\text{dep}}(\alpha)$ is retrieved from the ephemeris kernel via `cspice_spkezr` (object `-60000`), and the selenocentric rotation matrix $\mathbf{R}_{\text{sel}}(\alpha)$ is recomputed from the Moon’s state at the same epoch using the construction of Eq. (4.4). The pre-tabulated SPICE grid is shifted accordingly so that its origin aligns with $t_{\text{dep}}(\alpha)$.

The objective and penalty structure of Eqs. (4.14)–(4.15) remain unchanged in this extended formulation; only the initialisation of the forward propagation and the computation of the rotation matrix are affected. The bounds on α are set to the full normalised period $[0, 1)$, allowing the optimiser to select any departure point on the NRHO. Solutions with α close to α_{peri} are expected to converge to manoeuvre profiles compatible with the perilune-anchored solutions of Cases 1–4, while solutions at α values far from the perilune are expected to exhibit larger departure impulses (the spacecraft is farther from the Moon) but potentially shorter or more flexible descent geometries.

4.5 PSO and SQP Configuration

The optimisation pipeline employs a single PSO configuration for both scenarios, deliberately avoiding scenario- or case-specific tuning of the solver. This choice ensures that any variation of the converged ΔV across scenarios and across the cumulative constraint cases reflects the structure of the trajectory-optimisation problem itself and not the behaviour of the search algorithm. The same principle is extended to the local refinement stage: SQP

is applied uniformly to every case and every multi-start seed, with identical tolerances. The PSO is implemented via MATLAB’s `particleswarm` function (Global Optimisation Toolbox, R2024b) [41], building on the algorithm described in Section 3.3.1. The parameter values adopted are summarised in Table 4.5 and are common to all cases of both scenarios; the only run-to-run differences are the dimension of the decision vector (12 in the baseline cases, 13 when the NRHO phase α is admitted as a free variable) and the variable bounds of Tables 4.1 and 4.3.

Table 4.5. PSO solver parameters (`particleswarm`, MATLAB R2024b), common to both scenarios and to all cumulative constraint cases.

Parameter	MATLAB option	Value
Swarm size	<code>SwarmSize</code>	80
Maximum iterations	<code>MaxIterations</code>	500
Stall iterations	<code>MaxStallIterations</code>	500
Function tolerance	<code>FunctionTolerance</code>	0
Inertia range	<code>InertiaRange</code>	[0.4, 0.9]
Cognitive weight c_1	<code>SelfAdjustmentWeight</code>	2.0
Social weight c_2	<code>SocialAdjustmentWeight</code>	2.0
Independent seeds	<code>rng(k)</code> , $k = 1, \dots, N$	$N = 10$

The swarm size of 80 particles provides sufficient coverage of the 12- and 13-dimensional search spaces, as verified by preliminary sensitivity runs in which the swarm size was varied between 20 and 200 particles for both scenarios. The inertia weight is decreased linearly from 0.9 to 0.4 across iterations by MATLAB’s adaptive schedule [39], promoting broad exploration early and focused exploitation near convergence. The acceleration coefficients $c_1 = c_2 = 2.0$ are consistent with the original Kennedy–Eberhart formulation, which ensures mean-square convergence provided $\phi = c_1 + c_2 > 4$ [37]. The choices `MaxStallIterations` = `MaxIterations` and `FunctionTolerance` = 0 deliberately disable early stopping: each run is guaranteed to consume the full iteration budget, so that the comparison across cases and seeds is performed at a uniform amount of search effort rather than at variable, fitness-dependent termination points.

Convergence behaviour and choice of iteration count. The selection of 500 iterations was verified through a dedicated convergence study, in which the PSO was run on the most constrained case of Scenario 1 (Case 4a) over $N = 10$ independent seeds with deterministic initialisation `rng(k)`, $k = 1, \dots, 10$. For each seed the relative variation of the global best fitness over the last 100 iterations was computed as

$$\varepsilon_k = \frac{J_{N_{\text{it}}-100}^{(k)} - J_{N_{\text{it}}}^{(k)}}{J_{N_{\text{it}}}^{(k)}}, \quad N_{\text{it}} = 500. \quad (4.19)$$

The median value across the ten seeds is $\tilde{\varepsilon} = 0.77\%$ and the worst-case value is $\varepsilon_{\text{max}} = 2.28\%$; the full distribution of ε_k is reported in Table 4.6. The complete convergence histories are shown in Figure 4.6: nine of the ten seeds attain a final fitness in the range [3.85, 4.30] km/s, while a single outlier (seed 8) becomes trapped early in a higher-cost basin and stabilises around 6.73 km/s. The aggregated behaviour of the swarm is summarised in Figure 4.7,

which displays the median trajectory together with the interquartile band. Inspection of the individual tails over the last 150 iterations, shown in Figure 4.8, confirms that the residual descent observed at the end of every run is not due to undertraining but to PSO’s stochastic discovery of incremental improvements within each particle’s basin of attraction: the curves exhibit a stepwise pattern of discrete plateaus interrupted by sporadic step-downs of decreasing magnitude, rather than smooth monotonic decay. Extending the iteration budget beyond 500 would not eliminate these step-downs, which are an intrinsic feature of the multimodal landscape, while it would substantially increase the computational cost of the multi-start campaign described below. The chosen value 500 therefore represents a deliberate trade-off: it lies past the rapid descent phase of the search ($N_{\text{it}} \lesssim 100$) and well into the regime of diminishing returns, while remaining computationally sustainable when combined with multi-start.

Table 4.6. Relative variation of the PSO best fitness over the last 100 iterations, per seed, for the Case 4a convergence study ($N = 10$, $N_{\text{it}} = 500$).

Seed k	1	2	3	4	5	6	7	8	9	10	median	max
ε_k [%]	1.86	0.08	1.57	2.28	0.55	0.25	0.30	0.78	0.75	1.12	0.77	2.28




Figure 4.6. PSO convergence histories for Case 4a, $N = 10$ independent seeds with deterministic initialisation $\text{rng}(k)$. Nine seeds converge to final fitness values in the range $[3.85, 4.30]$ km/s, while seed 8 (grey) is trapped in a higher-cost basin and stabilises around 6.73 km/s, illustrating the multimodal character of the search space and motivating the multi-start strategy.




Figure 4.7. Median convergence trajectory and interquartile band (25th–75th percentile) across the $N = 10$ seeds of the Case 4a convergence study. The median plateaus past iteration ~ 200 and the IQR collapses to a narrow band around the median, indicating that the typical seed has reached its basin of attraction well before the end of the iteration budget.




Figure 4.8. Detail of the last 150 iterations of each seed for the Case 4a convergence study. The relative variation ε_k over the last 100 iterations is annotated on each panel. All seeds exhibit a stepwise descent with discrete plateaus interrupted by sporadic step-downs of decreasing magnitude, characteristic of PSO on multimodal landscapes, rather than the smooth monotonic decay typical of undertrained runs.

Multi-start strategy. The convergence study also reveals a substantial inter-seed dispersion of the converged fitness, with a mean of 4.33 km/s, a standard deviation of 0.86 km/s, and a corresponding coefficient of variation of 19.8% over the ten runs. This dispersion is not the by-product of incomplete convergence of the individual seeds, as established by the ε_k analysis above, but reflects the multimodality of the optimisation landscape: different seeds converge to qualitatively different basins of attraction, each

corresponding to a distinct trajectory geometry, and once trapped in a sub-optimal basin a seed cannot escape it within the iteration budget. The most extreme manifestation of this phenomenon is seed 8, which converges to $J = 6.73$ km/s, approximately 75% higher than the best seed of the same campaign; excluding this outlier the remaining nine seeds exhibit a coefficient of variation of 3.5%, comparable in magnitude to the per-seed convergence indicator ε_k . Crucially, increasing the PSO iteration budget would not remedy this inter-basin variability: a seed trapped in a worse basin at iteration 500 would remain trapped in the same basin at iteration 1000 or 5000, with only marginal further refinement of J within that basin.

The only structural remedy is a multi-start strategy, in which each case is solved with $N = 10$ independent PSO+SQP runs initialised with deterministic seeds `rng(1)` through `rng(10)`. The same set of seeds is reused across all cases and both scenarios, so that any inter-case variation of the reported ΔV is attributable to the change of the active constraint set rather than to a change of the initial random state. The case result is defined as the run that minimises the total impulsive cost $J_{1,2}$ among feasible solutions, that is, among solutions whose maximum nonlinear-constraint violation satisfies $\max_k c_k(\mathbf{x}_{\text{SQP}}^*) \leq 10^{-2}$. If, for a given case, no feasible solution is obtained from the ten runs, the least-infeasible run is reported with an explicit caveat.

SQP local refinement. Sequential Quadratic Programming is applied uniformly as the local refinement stage of every PSO run in the multi-start campaign. That is, each of the $N = 10$ PSO global bests $\mathbf{x}_{\text{PSO}}^{*,(k)}$, $k = 1, \dots, 10$, is warm-started into a separate SQP call, and the N refined solutions $\mathbf{x}_{\text{SQP}}^{*,(k)}$ enter the multi-start selection described above. The SQP refinement is carried out with MATLAB's `fmincon` configured with the '`sqp`' algorithm. The constraints are translated into native nonlinear inequalities $\mathbf{c}(\mathbf{x}) \leq \mathbf{0}$; the constraint, optimality and step tolerances are kept at their MATLAB default values (10^{-6} , 10^{-6} , 10^{-10} respectively). The objective minimised by `fmincon` is the unaugmented total impulsive cost $J_{1,2}$ of Eqs. (4.8) and (4.14), since the constraints are now enforced explicitly rather than through the penalty terms that augment the PSO fitness.

4.6 Optimisation Pipeline

The end-to-end workflow is structurally identical for both scenarios, differing only in three localised aspects: the direction of propagation (backward in mission time for Scenario 1, forward for Scenario 2, as motivated in Sections 4.1 and 4.4.1), the identification of the fixed NRHO state (apolune or perilune), and the form of the terminal event. For each case the workflow is executed $N = 10$ times with independent deterministic seeds, and the case result is selected according to the criterion of Section 4.5. The following five steps are performed at each multi-start invocation.

1. **Initialisation.** The SPICE kernels (`naif0012.tls`, `de440.bsp`, `receding_horiz_3189_1burnApo`) are loaded via `cspice_furnsh`. The fixed NRHO state is identified by a uniform period scan: $N_{\text{scan}} = 624$ samples (one every 15 min) over T_{NRHO} for Scenario 1 (apolune, Eq. (4.6)), or $N_{\text{scan}} = 3000$ samples for Scenario 2 (perilune, Eq. (4.11)). The corresponding frame rotation matrix - $\mathbf{R}$ in the lunar orbital plane frame (Eq. (4.2)) for Scenario 1, or $\mathbf{R}_{\text{sel}}$ in the selenocentric frame (Eq. (4.4)) for Scenario 2 - is computed once at initialisation and held fixed for the entire run, except when

the NRHO phase α is admitted as a decision variable, in which case the matrix is recomputed at every fitness evaluation. For Scenario 2 the Moon and Sun ephemerides are additionally pre-tabulated on a $\Delta t_{\text{grid}} = 120$ s grid over the maximum TOF window to accelerate the inner integration loop. At each of the $N = 10$ multi-start invocations the MATLAB random stream is reseeded via `rng(k)` with k the current seed index, ensuring reproducibility and consistency of the seed convention across all cases.

2. **PSO global search.** `particleswarm` evaluates the augmented cost $\tilde{J}_{1,2}(\mathbf{x})$ for the 80 particles of the swarm, updating positions and velocities according to Eqs. (3.5)–(3.6) for the full 500 iterations of the budget specified in Table 4.5.
3. **Trajectory propagation.** For each candidate decision vector, the propagation is built sequentially in the J2000 ECI frame, starting from the fixed NRHO state. The velocity vector is reversed for backward integration in Scenario 1 only. The first impulsive manoeuvre is applied via Eq. (4.1); segment 1 is integrated for τ_1 days. The second manoeuvre is then applied and segment 2 is integrated for τ_2 days. The third manoeuvre is applied and segment 3 is integrated for up to τ_3 days (with a $3\tau_3$ safety margin in Scenario 2 to accommodate delayed event triggering), under terminal event detection:
 - Scenario 1: ascending perigee, $\mathbf{r} \cdot \mathbf{v} = 0$;
 - Scenario 2: PDI altitude gate, $h(t) = 15$ km with $\dot{h}(t) < 0$, Eq. (4.12).
4. **Objective and constraint evaluation.** Terminal-state quantities are extracted (h_{peri} , i , t_{belt} , Δv_{circ} for Scenario 1; h_{PDI}^* , φ_{PDI} , τ_{tot} , h_{min} for Scenario 2) and the augmented cost $\tilde{J}_{1,2}$ of Eqs. (4.10) or (4.15) is returned to the PSO.
5. **SQP local refinement.** At the end of each PSO run, the global best $\mathbf{x}_{\text{PSO}}^*$ is passed to `fmincon` ('`sqp`') for local refinement under the explicit nonlinear-inequality formulation of Section 4.5, with default tolerances. The refined solution $\mathbf{x}_{\text{SQP}}^*$ and its maximum constraint-violation indicator are returned to the multi-start driver. Once all $N = 10$ runs are complete, the case result is selected as the run with the lowest ΔV among the feasible solutions, as defined in Section 4.5.

5 | Scenario 1: LEO $\rightarrow$ NRHO

This chapter presents the results of the hybrid PSO-SQP optimisation of the LEO $\rightarrow$ NRHO transfer scenario formulated in Chapter 4. The objective is to quantify how the minimum-cost transfer trajectory - and in particular its total ΔV budget - responds to the progressive introduction of operational constraints. As anticipated, rather than solving a single fully-constrained problem, the analysis follows a cumulative-constraint strategy: starting from a minimally-constrained baseline, each subsequent case retains all the constraints of the preceding one and activates exactly one additional operational requirement. This incremental structure isolates the marginal effect of each constraint on the cost and on the geometry of the optimal trajectory and exposes the mechanisms through which the optimiser redistributes the manoeuvre budget to accommodate the tightening feasible region.

All cases share the dynamical model, the three-impulse transcription, and the backward-propagation pipeline described in Chapter 4. Unless otherwise stated, the reference NRHO arrival epoch is the representative Artemis IV date introduced in Section 4.3.1, and the departure state is anchored at the apolune of the reference periodic orbit ($\alpha_{\text{apo}} = 0.8539$). Throughout the chapter, the “best” seed of each multi-start run is the feasible seed of lowest cost at the operational feasibility tolerance $\|\mathbf{c}\|_{\infty} \leq 10^{-3}$.

Roadmap of the constraint cases

The cumulative-constraint sequence comprises the following cases, summarised in Table 5.1.

- **Baseline (Case 1)** - Section 5.1. Only the perigee-altitude bound ($150 \text{ km} \leq h_p \leq 600 \text{ km}$) is enforced, which is the minimal constraint required for the LEO end-state to be physically well-posed. The time of flight is left free. The baseline provides an absolute lower bound on the transfer cost in the three-impulse parametrisation and the benchmark against which all subsequent marginal penalties are measured.
- **Case 2 (time of flight)** - the operational bound $\tau_{\text{tot}} \leq 5 \text{ d}$ is added. As discussed in Section 5.1, the baseline optimum already satisfies this bound with margin ($\tau_{\text{tot}} = 4.88 \text{ d}$), so the constraint is effectively non-binding at the global optimum and does not warrant a separate full analysis.
- **Case 3 (launch-site inclination)** - Sections 5.2 (Cape Canaveral, $26^\circ \leq i \leq 31^\circ$, Case 3a) and 5.3 (Kourou, $4^\circ \leq i \leq 8^\circ$, Case 3b). The LEO orbital plane is restricted to the inclination window achievable from each launch site. This is the first constraint

that is genuinely binding for both branches, and it introduces the large plane-change penalty that dominates the cost for the transfer from Kourou.

- **Case 4 (Van Allen radiation transit)** - Section 5.4. Separate upper bounds are imposed on the cumulative transit time through the inner ($t_{\text{inner}} \leq 0.50$ h) and outer ($t_{\text{outer}} \leq 2.00$ h) Van Allen belts, with thresholds derived from the Artemis I flight measurements [48]. The operational thresholds (Cases 4a/4b) are studied for both launch sites, together with a tightened-outer-belt sensitivity analysis ($t_{\text{outer}} \leq 1.50$ h, Cases 4c/4d).
- **Case 5 (free NRHO arrival point)** - Sections 5.5 (Cape Canaveral) and 5.7 (Kourou). The apolune-lock is released and the NRHO arrival phase α is admitted as a 13-th decision variable, allowing the optimiser to select the insertion point along the orbit. The marginal benefit of this additional degree of freedom is quantified for both launch sites, and its sensitivity to the arrival epoch is examined in the secular-variation study of Section 5.9.

Table 5.1. Roadmap of the cumulative-constraint cases for the LEO $\rightarrow$ NRHO scenario. "✓" marks the retained constraints and "+" marks the new ones.

Case	Launch site	h_p	$\tau_{\text{tot}} \leq 5$ d	Inclination	Van Allen / free α
1 (baseline)	—	+	free	—	—
2	—	✓	+	—	—
3a	Cape Canaveral	✓	✓	+	—
3b	Kourou	✓	✓	+	—
4a	Cape Canaveral	✓	✓	✓	+ Van Allen (op.)
4b	Kourou	✓	✓	✓	+ Van Allen (op.)
4c	Cape Canaveral	✓	✓	✓	Van Allen (tight.)
4d	Kourou	✓	✓	✓	Van Allen (tight.)
5a	Cape Canaveral	✓	✓	✓	✓ + free α
5b	Kourou	✓	✓	✓	✓ + free α

The remainder of this chapter analyses each case in turn, following a consistent structure: the optimal decision vector, the geometry of the optimised trajectory, the ΔV budget with comparison against the relevant baseline, and the multi-start convergence with verification of the active constraints. Section 5.10 collects the headline results across all cases and draws the chapter's overall conclusions.

5.1 Baseline Solution: Perigee Altitude Constraint Only

This section establishes the reference solution against which all subsequent results in this chapter are compared. The baseline scenario activates only the minimal constraint required for the optimisation problem to be physically well-posed: a perigee altitude bounded between 150 km and 600 km, which is the operational definition of the LEO end-state. All the other operational constraints are deliberately omitted. In particular the time of

flight is left as a free variable in Case 1; the upper bound $\tau_{\text{tot}} \leq 5$ d, which defines Case 2, is introduced only in the following case of the cumulative sequence. The baseline solution thus provides both an absolute lower bound on the transfer cost in the three-impulse parametrisation and a benchmark for quantifying the marginal ΔV penalty introduced by each subsequent operational requirement.

Throughout this chapter, the “best” seed of each multi-start optimisation run is the one that minimises the cost function $J(\mathbf{x})$ among those satisfying the operational feasibility tolerance $\|\mathbf{c}(\mathbf{x})\|_{\infty} \leq 10^{-3}$, where $\mathbf{c}$ collects all the problem constraints.

5.1.1 Optimal Decision Vector

The hybrid PSO–SQP scheme was executed with the bounds reported in Table 4.1 and the solver parameters of Table 4.5, over ten independent multi-start seeds. The best feasible seed converges to the solution summarised in Table 5.2, with a total mission cost $J^* = 3.783$ km/s, a total flight time $\tau_{\text{tot}} = 4.88$ d, and an LEO perigee at $h_p = 600$ km, $i = 40.6^\circ$. The remaining nine seeds are analysed in Section 5.1.4: as will be shown there, two distinct local-optimum basins emerge from the multi-start, with the global optimum reported here lying in the short-flight-time basin.

The most salient feature of the optimal decision vector is that both mid-course corrections collapse to negligible magnitudes: $\|\Delta \mathbf{v}_{\text{mid1}}\| = 1.8$ m/s and $\|\Delta \mathbf{v}_{\text{mid2}}\| = 0.4$ m/s, that is, between two and four orders of magnitude smaller than $\|\Delta \mathbf{v}_{\text{NRI}}\| = 740$ m/s. The three-impulse parametrisation thus effectively degenerates into a single-burn capture pattern: the spacecraft coasts ballistically from the LEO perigee to the NRHO apolune, and the entire on-board ΔV is concentrated in the arrival manoeuvre. The two intermediate nodes are retained by the optimiser only as residual numerical refinements. This is a strong, geometry-driven result: in the absence of operational constraints, the cheapest LEO-to-apolune transfer in the full-ephemeris model is essentially the direct apolune-tangent arc identified analytically by Zimovan et al. in the CR3BP [13], and the additional degrees of freedom granted by the three-impulse parametrisation are not exploited. The behaviour of the mid-course corrections will change qualitatively once operational constraints are activated in the subsequent sections; their inactivity in the baseline therefore quantifies the reserve of manoeuvring authority available to absorb future constraints.

Table 5.2. Optimal decision vector $\mathbf{x}^*$ - baseline scenario (LEO $\rightarrow$ NRHO, perigee altitude constraint only). Velocity components are expressed in the lunar orbital plane-aligned frame.

Variable	Description	Value	Unit
$\Delta v_{\text{NRI},x}$	NRHO insertion impulse, x	+0.0047	km/s
$\Delta v_{\text{NRI},y}$	NRHO insertion impulse, y	+0.7397	km/s
$\Delta v_{\text{NRI},z}$	NRHO insertion impulse, z	−0.0098	km/s
τ_1	Duration, segment 1	0.211	days
$\Delta v_{\text{mid1},x}$	Mid-course 1, x	−0.0004	km/s
$\Delta v_{\text{mid1},y}$	Mid-course 1, y	+0.0018	km/s
$\Delta v_{\text{mid1},z}$	Mid-course 1, z	+0.0000	km/s
τ_2	Duration, segment 2	0.763	days
$\Delta v_{\text{mid2},x}$	Mid-course 2, x	+0.0003	km/s
$\Delta v_{\text{mid2},y}$	Mid-course 2, y	+0.0002	km/s
$\Delta v_{\text{mid2},z}$	Mid-course 2, z	+0.0001	km/s
τ_3	Duration, segment 3 (realised)	3.909	days
<i>Derived quantities and diagnostics</i>			
$\ \Delta \mathbf{v}_{\text{NRI}}\ $	NRHO insertion magnitude	0.7398	km/s
$\ \Delta \mathbf{v}_{\text{mid1}}\ $	Mid-course correction 1 magnitude	0.0018	km/s
$\ \Delta \mathbf{v}_{\text{mid2}}\ $	Mid-course correction 2 magnitude	0.0004	km/s
J^*	Total mission cost	3.783	km/s
τ_{tot}	Total time of flight	4.88	days
h_p	LEO perigee	600	km
i	LEO inclination	40.6	°
α_{apo}	NRHO arrival phase (fixed)	0.8539	—
$d_{\text{Moon,min}}$	Minimum spacecraft–Moon distance	70 988	km

5.1.2 Trajectory Geometry

The optimised trajectory is displayed in the lunar orbital plane-aligned ECI frame in Figure 5.1, with a detailed view of the segments near the NRHO arrival point provided in Figure 5.2. The corresponding selenocentric representation is shown in Figure 5.3. Three features of the optimised transfer deserve attention.

- The transfer is direct, with no lunar flyby: the minimum spacecraft–Moon distance along the arc is 70 988 km, approximately 5 000 km outside the lunar sphere of influence ($\sim 66\,000$ km). The margin is modest but strictly positive, and the geometry is consistent with the class of low-sensitivity direct transfers identified by Zimovan et al. as the most navigationally robust option for crewed access to the L_2 NRHOs [13].
- The arrival point on the NRHO is fixed by construction at the apolune passage identified by the scan procedure of Eq. (4.6). For the reference epoch of 10 March 2028, 00:00 UTC, this scan returns a normalised NRHO phase $\alpha_{\text{apo}} = 0.8539$.
- The three segment durations are highly asymmetric, with $\tau_1 = 0.21$ d, $\tau_2 = 0.76$ d, and $\tau_3 = 3.91$ d. The two mid-course nodes therefore cluster within the first day of backward flight, that is, within the first day preceding the NRHO insertion in forward time. The remaining $\sim 80\%$ of the transfer is a long ballistic coast on

segment 3 down to the LEO perigee. Combined with the near-zero magnitudes of $\Delta \mathbf{v}_{\text{mid1}}$ and $\Delta \mathbf{v}_{\text{mid2}}$, this confirms that the mid-course nodes are geometrically present but kinematically inactive.

LEO → NRHO transfer (3 impulses) - Case 1 Baseline - Lunar-plane-aligned (rotated J2000)




Figure 5.1. Baseline optimal transfer in the lunar orbital plane-aligned ECI frame.




Figure 5.2. Close-up of the NRHO arrival region in the lunar orbital plane-aligned ECI frame.




Figure 5.3. Baseline optimal transfer in the selenocentric frame. The lunar sphere of influence (SOI, radius 66 000 km) is shown as a faint grid.

5.1.3 ΔV Budget

The total cost of the optimal baseline transfer is decomposed in Table 5.3. It is convenient to distinguish two contributions: the trans-lunar injection (TLI) cost Δv_{circ} , which is typically provided by the launch vehicle upper stage and is therefore not borne by the spacecraft propulsion system, and the on-board cost $\Delta V_{\text{sc}} = \|\Delta \mathbf{v}_{\text{NRI}}\| + \|\Delta \mathbf{v}_{\text{mid1}}\| + \|\Delta \mathbf{v}_{\text{mid2}}\|$, which is the relevant metric for sizing the spacecraft propulsion system. Both contributions are nonetheless included in the objective function (4.8), so that the optimiser can trade them against each other.

Table 5.3. ΔV budget – baseline scenario.

Manoeuvre	Δv [km/s]	Share of J^* [%]
LEO insertion (TLI, Δv_{circ})	3.041	80.4
Mid-course correction 1	0.002	0.0
Mid-course correction 2	0.000	0.0
NRHO insertion ($\Delta \mathbf{v}_{\text{NRI}}$)	0.740	19.6
On-board total ΔV_{sc}	0.742	19.6
Total J^*	3.783	100.0

Two observations are appropriate. First, ΔV_{sc} is, to within 0.3%, equal to $\|\Delta \mathbf{v}_{\text{NRI}}\|$ alone: the optimised on-board budget is borne almost entirely by the single capture burn at the NRHO apolune, and the spacecraft propulsion system can in principle be sized from this manoeuvre alone. This is the confirmation of the degeneracy already noted in Section 5.1.1.

Second, the optimised total cost of 3.783 km/s lies modestly below the 3.953 km/s reported by Zimovan et al. for a comparable direct apolune-tangent transfer from a circular LEO at 28° inclination [13]. The difference of ~ 170 m/s is primarily attributable to the higher perigee altitude allowed by the present formulation (600 km, at the upper bound, versus the 200 km of the reference work), which lowers the required circular velocity at perigee by approximately 200 m/s and is therefore consistent with the observed reduction. The marginal contribution from the two mid-course corrections, which the reference CR3BP transfer does not include, is here negligible.

A complementary view of the dominant manoeuvre at the LEO end is shown in Figure 5.4, drawn in the pure ECI J2000 frame (rather than in the lunar orbital plane-aligned frame used by all previous figures of this section) so that the parking orbit inclination is directly visible against the Earth equatorial plane. The implicit TLI impulse Δv_{circ} is essentially tangential to the parking orbit at perigee, consistently with the classical result that a tangential burn is optimal for energy injection.




Figure 5.4. Baseline LEO departure region in the ECI J2000 frame.

5.1.4 Multi-start Convergence, Active Constraints, and Operational Look-ahead

Multi-start convergence and the structure of the solution space. The hybrid PSO–SQP scheme was run with ten independent random seeds. All ten seeds converged to operationally feasible solutions ($\|\mathbf{c}\|_\infty \leq 10^{-3}$); the summary of all ten outcomes, sorted by total mission cost, is reported in Table 5.4. Two distinct local-optimum basins emerge clearly from the multi-start outcome: a short-flight-time basin populated by two seeds (Seeds 7 and 2, with $\tau_{\text{tot}} \in [4.88, 5.05]$ d), and a long-flight-time basin populated by the remaining eight seeds (with $\tau_{\text{tot}} \in [15.4, 17.0]$ d). The global optimum lies in the short-flight-time basin (Seed 7, $J^* = 3.783$ km/s); the best long-flight-time solution (Seed 3, $J = 3.800$ km/s) is only 17 m/s (0.45%) above it.

Table 5.4. Multi-start convergence of the baseline scenario.

Seed	J [km/s]	τ_{tot} [d]	$\ \Delta \mathbf{v}_{\text{NRI}}\ $ [km/s]	h_p [km]	i [deg]	Basin
7	3.783	4.88	0.740	600.0	40.6	short
2	3.797	5.05	0.421	600.0	38.2	short
3	3.800	15.41	0.760	599.6	37.4	long
5	3.805	15.81	0.635	600.0	39.5	long
8	3.813	16.39	0.246	600.0	35.3	long
6	3.817	16.97	0.445	600.0	38.8	long
10	3.833	16.94	0.449	600.0	40.7	long
1	3.857	15.41	0.789	600.0	42.1	long
4	3.861	15.52	0.806	600.0	40.1	long
9	3.947	16.70	0.605	600.0	46.0	long

The optimiser’s modest preference for the short-flight-time basin can be attributed to two complementary mechanisms.

- **Perturbation accumulation along the coast.** The long-flight-time trajectories spend two to three times longer under the action of the lunar and solar third-body perturbations of Eq. (2.2). Even when no operational constraint is enforced, the optimiser must compensate the resulting position and velocity drift at the NRHO insertion, which is paid in additional $\Delta \mathbf{v}_{\text{NRI}}$. This effect is visible in several long-flight-time seeds whose $\|\Delta \mathbf{v}_{\text{NRI}}\|$ exceeds the short-flight-time optimum (0.760 km/s for Seed 3, 0.789 km/s for Seed 1, 0.806 km/s for Seed 4, against 0.740 km/s of Seed 7).
- **Distinct geometric class of transfer.** The two basins do not differ by a mere stretching of the same trajectory: they correspond to qualitatively different arc topologies. The short-flight-time basin features a single long ballistic coast on segment 3 with the mid-course nodes packed near the NRHO end, while the long-flight-time basin features two long coasts with the mid-course nodes distributed along the cislunar leg. The same parametric freedom that yields no advantage here will become genuinely advantageous in Case 3, once the LEO inclination constraint forces an out-of-plane correction.

The long-flight-time basin is not a low-energy transfer. It is worth emphasising that the long-flight-time solutions found above are not the low-energy or weak stability boundary (WSB) transfers used for cargo deliveries to the lunar vicinity. True WSB transfers exploit weak stability boundary dynamics in the Earth-Moon-Sun four-body problem to achieve ballistic capture into lunar orbits at ΔV costs of the order of 50–200 m/s below the direct value, in exchange for transfer durations of 70–110 days. The 15–17 day solutions reported in Table 5.4 are simply longer direct transfers, and offer no comparable ΔV advantage. They are therefore of no operational interest for crewed missions, where the time-of-flight upper bound of Case 2 will exclude them by construction.

Time of flight: a non-binding constraint of Case 2. The conjecture that the time-of-flight upper bound $\tau_{\text{tot}} \leq 5$ d of Case 2 is non-binding can now be tested directly: the same hybrid PSO–SQP scheme has been rerun with the Case 2 constraint activated

and otherwise identical multi-start configuration. The outcome is summarised in Table 5.5. Two observations support the non-binding claim quantitatively.

First, the global optimum of Case 2 coincides with that of Case 1 to within numerical tolerance. The best Case 2 seed converges to $J^* = 3.7819$ km/s at $\tau_{\text{tot}} = 4.86$ d, a difference of -1.0 m/s in cost and -0.02 d in flight time with respect to Case 1, both well within the operational feasibility tolerance of the solver. The mid-course corrections remain at the same negligible level ($\|\Delta \mathbf{v}_{\text{mid1}}\| = 0.2$ m/s and $\|\Delta \mathbf{v}_{\text{mid2}}\| = 0.6$ m/s) and the on-board cost is $\Delta V_{\text{sc}} = 741$ m/s, indistinguishable from the 742 m/s of Case 1. Trajectory geometry, ΔV budget, and constraint pattern of the Case 2 optimum are therefore identical to those of the Case 1 optimum reported in Sections 5.1.1–5.1.3, and a dedicated section for Case 2 is not justified.

Second, and most importantly for the cumulative-constraint argument of this chapter, none of the ten Case 2 multi-start seeds converged to the long-flight-time basin. All ten seeds returned $\tau_{\text{tot}} \in [4.45, 5.00]$ d, with the constraint mildly active (saturated at $\tau_{\text{tot}} = 5.0000$ d) only for two seeds. The TOF constraint of Case 2 therefore performs exactly the intended function of operationally excluding crew-incompatible long-coast trajectories from the feasible domain, while leaving the short-flight-time basin, which already hosts the Case 1 global optimum, substantially untouched.

Table 5.5. Effect of activating the Case 2 time-of-flight constraint on the multi-start optimisation outcome.

Scenario	Seed	J^* [km/s]	τ_{tot} [d]	$\ \Delta \mathbf{v}_{\text{NRI}}\ $ [m/s]	ΔV_{sc} [m/s]	long-TOF seeds
Case 1 best	7	3.7829	4.88	740	742	8/10
Case 2 best	1	3.7819	4.86	740	741	0/10

Active constraints at convergence. Restricting attention to the global optimum (Seed 7), the two active constraints of Case 1 are satisfied at convergence:

- **Perigee altitude:** $h_p = 600$ km, exactly at the upper bound of the admissible range $[150, 600]$ km. As Table 5.4 shows, this saturation is universal across the ten seeds, regardless of the basin. The optimiser drives the perigee towards the highest admissible altitude in order to minimise the TLI cost, and this saturation recurs in all subsequent cases of this chapter.
- **NRHO insertion:** the position and velocity at the end of segment 1 match the reference NRHO state at $\alpha_{\text{apo}} = 0.8539$ to within the operational feasibility tolerance $\|\mathbf{c}\|_{\infty} \leq 10^{-3}$.

Operational constraint look-ahead. For completeness, the global optimum is also evaluated against the operational constraints that will be activated in the subsequent cases of this chapter, even though they are not enforced in Case 1:

- **LEO inclination:** $i = 40.6^\circ$, outside both the Cape Canaveral (26° – 31° , Case 3a) and the Kourou (4° – 8° , Case 3b) admissible ranges. The unconstrained optimiser selects a transfer plane that no realistic launcher could deliver without an additional plane-change manoeuvre.

- **Van Allen radiation belt transit:** $t_{\text{inner}} = 0.32$ h within the inner belt and $t_{\text{outer}} = 1.73$ h within the outer belt, both within their respective operational thresholds of 0.50 h and 2.00 h that will be enforced in Case 4. The baseline trajectory is therefore already compatible with the Case 4 radiation-transit constraint; the relevance of that constraint in the subsequent cases of this chapter will be re-examined when the inclination requirements modify the transfer geometry.

The two violations above are not penalised in the baseline objective, but their magnitudes already signal the directions along which the optimiser will be forced to deform the trajectory once the corresponding constraints are activated. In particular, the $\sim 10^\circ$ misalignment between the unconstrained transfer plane and the Cape Canaveral inclination window is expected to be absorbed predominantly by the mid-course corrections, which the baseline shows are currently unused and therefore offer the largest reserve of in-plane manoeuvring authority; the analysis of this redistribution is the subject of the next section.

Sensitivity: tightening the time-of-flight bound. To confirm that the non-binding behaviour of Case 2 is a property of the 5-day threshold and not an artefact of the optimiser, a sensitivity run was performed with the bound tightened to $\tau_{\text{tot}} \leq 3.5$ d, all other constraints unchanged. Under this tighter bound the constraint becomes genuinely active: the best strictly-feasible seed converges to $J^* = 3.928 \text{ km s}^{-1}$ at $\tau_{\text{tot}} = 3.496$ d, a 146 m s^{-1} (+20% with respect to the on-board ΔV) increase over the Case 2 optimum, driven by a less favourably-aligned trans-lunar arc that must be compressed into a shorter flight time. Table 5.6 summarises the comparison.

Table 5.6. Sensitivity of the optimal cost to the time-of-flight bound.

Bound $\tau_{\text{tot}} \leq \cdot$	J^* [km/s]	τ_{tot}^* [d]	Binding?
5.0 d (Case 2)	3.782	4.88	No
3.5 d (sensitivity)	3.928	3.50	Yes

The result confirms the interpretation of Section 5.1: minimum-cost transfers in this class are naturally short-duration, so a generous TOF ceiling costs nothing, while a sufficiently aggressive one forces a less efficient trans-lunar geometry and is paid for in ΔV . A full mapping of the cost-versus-TOF trade-off across the 3.5-5.0 d range is left as a natural extension of this sensitivity check.

5.2 Effect of LEO Inclination Constraint: Cape Canaveral (Case 3a)

The first operational requirement layered on top of the baseline is the launch-site inclination constraint, which restricts the LEO orbital plane to the window achievable from a specified ground facility. The baseline solution of Section 5.1 converged to $i = 40.6^\circ$, outside both the Cape Canaveral and the Kourou admissible ranges; the inclination constraint is therefore expected to be binding in both Case 3a and Case 3b. The present section analyses the Cape Canaveral case ($26^\circ \leq i \leq 31^\circ$) in detail and quantifies the marginal

ΔV penalty with respect to the baseline. The corresponding Kourou case ($4^\circ \leq i \leq 8^\circ$) is addressed in Section 5.3, and the two launch sites are compared at the end of that section.

5.2.1 Optimal Decision Vector

The hybrid PSO–SQP scheme was executed with the same multi-start configuration of Case 1 and the Cape Canaveral inclination constraint $26^\circ \leq i \leq 31^\circ$ activated. The best operationally-feasible seed converges to the solution summarised in Table 5.7, with a total mission cost $J^* = 3.829$ km/s, a total flight time $\tau_{\text{tot}} = 5.00$ d (at the Case 2 upper bound), and an LEO perigee at $h_p = 600$ km, $i = 31.0^\circ$.

The most salient feature of the optimal decision vector is the exploitation of the mid-course node 1, which had collapsed to a negligible magnitude in the baseline and now carries $\|\Delta \mathbf{v}_{\text{mid1}}\| = 318$ m/s, comparable in size to the NRHO insertion impulse itself. The second mid-course remains essentially inactive ($\|\Delta \mathbf{v}_{\text{mid2}}\| = 0.6$ m/s), and the three-impulse parametrisation operates effectively as a two-burn scheme: a substantial plane-and-energy manoeuvre at mid1, followed by a reduced capture impulse at NRI. The detailed structure of $\Delta \mathbf{v}_{\text{mid1}}$, its decomposition in the local RTN frame, and the physical motivation for its location along the transfer arc are discussed in Section 5.2.3.

A further structural change with respect to the baseline concerns the segment durations. Whereas in Case 1 the two mid-course nodes were clustered within the first day of backward flight ($\tau_1 = 0.21$ d, $\tau_2 = 0.76$ d), in Case 3a they are spread across the transfer ($\tau_1 = 1.46$ d, $\tau_2 = 2.46$ d), so that the long ballistic coast on segment 3 collapses from 3.91 d to 1.07 d. The first node (mid1) is now located near the end of the cislunar coast in forward time, where the spacecraft velocity is substantially lower than at perigee, and the second node (mid2) acts as a residual correction just before the LEO arrival.

Table 5.7. Optimal decision vector $\mathbf{x}^*$ - Case 3a (Cape Canaveral inclination constraint $26^\circ \leq i \leq 31^\circ$). Velocity components are expressed in the lunar orbital plane-aligned frame.

Variable	Description	Value	Unit
$\Delta v_{\text{NRI},x}$	NRHO insertion impulse, x	-0.0035	km/s
$\Delta v_{\text{NRI},y}$	NRHO insertion impulse, y	$+0.4700$	km/s
$\Delta v_{\text{NRI},z}$	NRHO insertion impulse, z	$+0.0141$	km/s
τ_1	Duration, segment 1	1.464	days
$\Delta v_{\text{mid1},x}$	Mid-course 1, x	$+0.0145$	km/s
$\Delta v_{\text{mid1},y}$	Mid-course 1, y	$+0.3000$	km/s
$\Delta v_{\text{mid1},z}$	Mid-course 1, z	-0.1033	km/s
τ_2	Duration, segment 2	2.465	days
$\Delta v_{\text{mid2},x}$	Mid-course 2, x	$+0.0000$	km/s
$\Delta v_{\text{mid2},y}$	Mid-course 2, y	$+0.0003$	km/s
$\Delta v_{\text{mid2},z}$	Mid-course 2, z	-0.0005	km/s
τ_3	Duration, segment 3 (realised)	1.302	days
<i>Derived quantities and diagnostics</i>			
$\ \Delta \mathbf{v}_{\text{NRI}}\ $	NRHO insertion magnitude	0.4702	km/s
$\ \Delta \mathbf{v}_{\text{mid1}}\ $	Mid-course correction 1 magnitude	0.3176	km/s
$\ \Delta \mathbf{v}_{\text{mid2}}\ $	Mid-course correction 2 magnitude	0.0006	km/s
J^*	Total mission cost	3.829	km/s
τ_{tot}	Total time of flight	5.000	days
h_p	LEO perigee	600	km
i	LEO inclination	31	$^\circ$
α_{apo}	NRHO arrival phase (fixed)	0.8539	—
$d_{\text{Moon,min}}$	Minimum spacecraft–Moon distance	$70\,932$	km

5.2.2 Trajectory Geometry

The optimised trajectory is displayed in the lunar orbital plane-aligned ECI frame in Figure 5.5, and in the selenocentric frame in Figure 5.6. Compared with the baseline, three features deserve attention.

- The transfer remains direct, with no lunar flyby: the minimum spacecraft–Moon distance is 70 932 km, still strictly outside the lunar sphere of influence. The constraint does not destabilise this property.
- The spatial distribution of the manoeuvre nodes is different from the baseline. The first mid-course node now sits visibly before the arrival, with a yellow arrow representing $\Delta \mathbf{v}_{\text{mid1}} = 318$ m/s. In the baseline this node was packed against the NRHO end.
- The NRHO insertion impulse is visibly smaller than in the baseline: the cyan arrow at the white star is approximately 36% times shorter, consistent with the reduction $740 \rightarrow 470$ m/s.

LEO → NRHO transfer (3 impulses) - Case 3a (Cape Canaveral $26^\circ < i < 31^\circ$) - Lunar-plane-aligned (rotated J2000)




Figure 5.5. Case 3a optimal transfer in the lunar orbital plane-aligned ECI frame, with the Cape Canaveral inclination constraint ($26^\circ \leq i \leq 31^\circ$) active.




Figure 5.6. Case 3a optimal transfer in the selenocentric frame.

5.2.3 ΔV Budget: Where the Inclination Penalty Is Paid

The total cost of the constrained solution is $J^* = 3.829$ km/s, a +46 m/s increase (+6.2% on the on-board ΔV) with respect to the baseline. The decomposition of this penalty is reported in Table 5.8 alongside the baseline values: the aggregate +46 m/s figure conceals a much more interesting redistribution of the on-board manoeuvre budget, which is the central result of this section.

Table 5.8. ΔV budget for Case 3a, compared with the Case 1 baseline. The total cost increases by only +46 m/s, but this aggregate figure conceals a substantial internal redistribution: +316 m/s migrate to the mid-course node 1, while -270 m/s are recovered at the NRHO insertion. The TLI cost is essentially unchanged.

Manoeuvre	Case 1 [m/s]	Case 3a [m/s]	Δ [m/s]
LEO insertion (TLI, Δv_{circ})	3041	3040	-1
Mid-course correction 1	2	318	+316
Mid-course correction 2	0	1	+1
NRHO insertion ($\Delta \mathbf{v}_{\text{NRI}}$)	740	470	-270
On-board total ΔV_{sc}	742	789	+47
Total J^*	3783	3829	+46

The physical interpretation of this redistribution is the following. In the baseline, the LEO inclination of $i = 40.6^\circ$ was chosen freely by the optimiser to align the LEO plane with the geometrically optimal transfer plane connecting the perigee to the NRHO apolune; the resulting trajectory required almost no in-flight plane change, and the on-board ΔV collapsed entirely onto the single capture burn at NRI. In Case 3a, the inclination is forced to $i \leq 31^\circ$ - the constraint saturates at the upper bound, $i = 30.999^\circ$ - and the LEO orbital plane is therefore misaligned from the optimal transfer plane by approximately 10° . This residual out-of-plane component must be removed somewhere along the trajectory, and the question of where is the key design decision the optimiser is making.

The choice settled upon by the optimiser is to perform the plane change at mid-course node 1, which is located at 1.46 d after the NRHO insertion in backward propagation - equivalently, at $\tau_{\text{tot}} - 1.46 = 3.54$ d after the TLI in forward time. At this point along the transfer the spacecraft has already coasted past $\sim 80\%$ of the cislunar leg in forward time and its inertial velocity has dropped by approximately one order of magnitude from its perigee value (from ~ 10.5 km/s at TLI to ~ 0.8 km/s at mid1). A plane-change manoeuvre at low velocity is energetically cheap: under the formula

$$\Delta v_{\text{pc}} = 2v \sin(\Delta i/2) \quad (5.1)$$

a 10° rotation at $v = 0.8$ km/s costs only ~ 140 m/s, compared with ~ 1.8 km/s if the same rotation were attempted at the TLI perigee. The factor-of-ten penalty avoided by performing the plane change at mid1 is precisely what makes the constrained transfer affordable in spite of the binding inclination requirement.

This interpretation is corroborated by the decomposition of the four impulses in the local Radial-Tangential-Normal (RTN) frame, computed at the pre-impulse spacecraft state and reported in Figure 5.7. Three observations follow.

- The TLI impulse is essentially purely tangential, as expected for an optimal energy-injection burn at perigee: any radial or normal component would introduce a parasitic plane mismatch that the optimiser systematically avoids.
- The NRI impulse is predominantly tangential as well, with a smaller normal contribution - consistent with its primary role as a braking manoeuvre at the NRHO apolune where the spacecraft velocity must be aligned with the NRHO tangent vector.
- The mid-course node 1 impulse exhibits a comparable contribution from the tangential and the normal directions, with a negligible radial component. This is the signature

of a combined manoeuvre in which a single burn simultaneously performs (i) a plane rotation (the normal component) and (ii) a residual energy correction (the tangential component) required because the constrained-plane trajectory has slightly different energy needs at mid1 than the unconstrained direct transfer would have. The mid-course node 2, in contrast, retains its baseline near-zero magnitude.




Figure 5.7. Decomposition of the four impulses of Case 3a in the local Radial-Tangential-Normal (RTN) frame, computed at the pre-impulse spacecraft state.

A complementary visualisation of the LEO end-state, drawn in the pure ECI J2000 frame to render the equatorial inclination visible against the equatorial plane, is provided in Figure 5.8. The orbital inclination of $i = 31.0^\circ$ saturates the upper bound of the Cape Canaveral admissible range and lies $\sim 9.6^\circ$ below the unconstrained baseline value of 40.6° .




Figure 5.8. Case 3a LEO departure region in the ECI J2000 frame.

From a mission-design perspective, the relevant quantity is the on-board budget $\Delta V_{sc} = 789$ m/s, which is borne by the spacecraft propulsion system. The Cape Canaveral inclination constraint inflates this quantity by +47 m/s (+6.3%) with respect to the baseline - a small absolute margin that nonetheless represents the cumulative cost of the operational requirement. The total figure J^* , which also includes the TLI provided by the launch vehicle upper stage, increases by only +1.2%.

5.2.4 Multi-start Convergence, Active Constraints, and Operational Look-ahead

Multi-start convergence. The ten independent multi-start seeds converged to a much narrower distribution than in Case 1: all ten returned $\tau_{tot} \in [4.35, 5.00]$ d (the long-flight-time basin is excluded by the Case 2 TOF constraint) and $i \in [27.7^\circ, 31.0^\circ]$ (within the Cape Canaveral admissible range, with a strong clustering at the 31° upper bound), as summarised in Table 5.9. Seven of the ten seeds converged within 50 m/s of the global optimum ($J \in [3.829, 3.872]$ km/s); the three outliers (Seeds 1, 3, 8) reached significantly higher costs (3.918–4.026 km/s) by saturating mid-course node 1 against its bound ($\|\Delta \mathbf{v}_{mid1}\| = 0.300$ km/s in Seeds 10, 6, 4, 5). This bound saturation is itself a diagnostic of the binding nature of the inclination constraint: the optimiser is pushing $\|\Delta \mathbf{v}_{mid1}\|$ as close to its allowed limit as the problem admits.

Table 5.9. Multi-start convergence of Case 3a (Cape Canaveral inclination constraint).

Seed	J [km/s]	τ_{tot} [d]	$\ \Delta \mathbf{v}_{\text{NRI}}\ $ [m/s]	$\ \Delta \mathbf{v}_{\text{mid1}}\ $ [m/s]	i [deg]
10	3.829	5.00	470	318	31.0
6	3.829	5.00	512	268	31.0
9	3.833	5.00	571	39	31.0
5	3.833	5.00	649	130	31.0
4	3.837	5.00	663	112	31.0
7	3.848	4.97	686	20	31.0
2	3.862	4.86	625	104	30.6
1	3.918	4.75	769	17	31.0
3	3.922	4.96	641	35	27.7
8	4.026	4.36	526	366	30.6

Active constraints at convergence. Restricting attention to the global optimum (Seed 10), three active constraints of Case 3a are satisfied at convergence:

- **Perigee altitude:** $h_p = 600$ km, exactly at the upper bound of the admissible range $[150, 600]$ km. The saturation, already universal in Case 1, is maintained in Case 3a.
- **Total time of flight:** $\tau_{\text{tot}} = 5.000$ d, exactly at the upper bound. Unlike Case 1–2, where the TOF constraint was operationally non-binding, in Case 3a the TOF constraint becomes marginally active: the inclination requirement forces a longer cislunar coast on segment 2 to allow the plane-change manoeuvre at mid1 to be performed at low velocity, and the optimiser settles against the TOF upper bound. This is the first qualitative change in the role of the TOF constraint along the cumulative-constraint sequence.
- **LEO inclination:** $i = 31.0^\circ$, saturating the upper bound of the Cape Canaveral range $[26^\circ, 31^\circ]$.
- **NRHO insertion equality:** the position and velocity at the end of segment 1 match the reference NRHO state at $\alpha_{\text{apo}} = 0.8539$ to within the operational feasibility tolerance $\|\mathbf{c}\|_\infty \leq 10^{-3}$.

Operational constraint look-ahead. For completeness, the global optimum is also evaluated against the Van Allen belts operational constraint that will be activated in the subsequent cases of this chapter. The Van Allen radiation belt transit times are $t_{\text{inner}} = 0.32$ h within the inner belt and $t_{\text{outer}} = 1.75$ h within the outer belt, both within their respective 0.50 h and 2.00 h operational thresholds that will be enforced in Case 4. The inclination constraint has therefore not pushed the Van Allen transit beyond its limits, consistent with the fact that radiation transit is primarily set by the energy of the trans-lunar arc rather than by its inclination: the outer-belt transit increases by only 0.02 h with respect to the baseline.

5.3 Effect of LEO Inclination Constraint: Kourou (Case 3b)

The Cape Canaveral case of Section 5.2 restricted the LEO inclination to a window that lies $\sim 10^\circ$ below the unconstrained baseline value of 40.6° , and the resulting marginal penalty on the on-board budget was modest (+47 m/s on ΔV_{sc}). The Kourou case considered here imposes a much tighter constraint, $4^\circ \leq i \leq 8^\circ$, which forces the LEO plane to lie $\sim 33^\circ$ below the unconstrained value. On the basis of the linearised plane-change cost analysis of Section 5.2.3, the penalty is expected to scale roughly with $\sin(\Delta i/2)$, that is, a factor of approximately 3.3 with respect to the Cape Canaveral case. The actual outcome is in fact considerably more severe than this linear estimate would suggest, and reveals a structural change in the optimal manoeuvre architecture that is the main subject of this section. The closing subsection 5.3.5 provides a quantitative side-by-side comparison of the two launch sites that is the natural concluding piece of this chapter’s inclination-constraint analysis.

5.3.1 Optimal Decision Vector

The hybrid PSO–SQP scheme was executed with the same multi-start configuration of Case 3a and the Kourou inclination constraint $4^\circ \leq i \leq 8^\circ$ activated. The best operationally-feasible seed converges to the solution summarised in Table 5.10, with a total mission cost $J^* = 4.261$ km/s, a total flight time $\tau_{tot} = 5.00$ d, and an LEO perigee at $h_p = 600$ km, $i = 8.0^\circ$.

Two qualitative features distinguish this decision vector from those analysed in the previous cases.

First, the second mid-course correction is fully exploited. Whereas $\|\Delta \mathbf{v}_{mid2}\|$ remained below 1 m/s in Cases 1, 2 and 3a, in Case 3b it reaches 347 m/s, of the same order as both $\|\Delta \mathbf{v}_{mid1}\| = 428$ m/s and $\|\Delta \mathbf{v}_{NRI}\| = 473$ m/s. Every degree of freedom is now exploited: all three on-board manoeuvres carry comparable shares of ΔV_{sc} .

Second, the duration of the segment between mid1 and mid2 has collapsed to zero: $\tau_2 = 0.000$ d to the precision of the solver, meaning that the two mid-course nodes are co-located in time along the transfer arc. The physical meaning of this co-location is discussed in detail in Section 5.3.3; in short, the two mid-course impulses act as a single compound manoeuvre of magnitude $\|\Delta \mathbf{v}_{mid1} + \Delta \mathbf{v}_{mid2}\| = 775$ m/s, split into two distinct impulse nodes only because the per-component bound $|\Delta v_{mid,k}| \leq 0.300$ km/s of the optimisation problem prevents a single impulse from carrying it in full. This split is mathematically non-unique: the multi-start analysis of Section 5.3.4 will show that Seeds 9 and 7 converge to identical mission cost and identical compound mid-course vector, but with two different individual splits of that vector between mid1 and mid2 – direct evidence that the parametrisation degenerates into the compound vector at this convergence point.

Table 5.10. Optimal decision vector $\mathbf{x}^*$ - Case 3b (Kourou inclination constraint $4^\circ \leq i \leq 8^\circ$). Velocity components are expressed in the lunar orbital plane-aligned frame. The duration of segment 3 is event-terminated by the perigee crossing detector.

Variable	Description	Value	Unit
$\Delta v_{\text{NRI},x}$	NRHO insertion impulse, x	−0.0950	km/s
$\Delta v_{\text{NRI},y}$	NRHO insertion impulse, y	+0.4290	km/s
$\Delta v_{\text{NRI},z}$	NRHO insertion impulse, z	+0.1745	km/s
τ_1	Duration, segment 1	3.755	days
$\Delta v_{\text{mid1},x}$	Mid-course 1, x	+0.2178	km/s
$\Delta v_{\text{mid1},y}$	Mid-course 1, y	+0.2455	km/s
$\Delta v_{\text{mid1},z}$	Mid-course 1, z	−0.2755	km/s
τ_2	Duration, segment 2	0.000	days
$\Delta v_{\text{mid2},x}$	Mid-course 2, x	+0.1762	km/s
$\Delta v_{\text{mid2},y}$	Mid-course 2, y	+0.1987	km/s
$\Delta v_{\text{mid2},z}$	Mid-course 2, z	−0.2229	km/s
τ_3	Duration, segment 3 (realised)	1.245	days
<i>Derived quantities and diagnostics</i>			
$\ \Delta \mathbf{v}_{\text{NRI}}\ $	NRHO insertion magnitude	0.4728	km/s
$\ \Delta \mathbf{v}_{\text{mid1}}\ $	Mid-course 1 magnitude	0.4285	km/s
$\ \Delta \mathbf{v}_{\text{mid2}}\ $	Mid-course 2 magnitude	0.3467	km/s
$\ \Delta \mathbf{v}_{\text{mid1}} + \Delta \mathbf{v}_{\text{mid2}}\ $	Compound mid magnitude	0.775	km/s
J^*	Total mission cost	4.261	km/s
τ_{tot}	Total time of flight	5.000	days
h_p	LEO perigee	600	km
i	LEO inclination	8	°
α_{apo}	NRHO arrival phase (fixed)	0.8539	—
$d_{\text{Moon,min}}$	Minimum spacecraft–Moon distance	69 988	km

5.3.2 Trajectory Geometry

The optimised trajectory is displayed in the lunar orbital plane-aligned ECI frame in Figure 5.9 and in the Moon-Centred Inertial (MCI) frame in Figure 5.10. The qualitative comparison with the previous cases reveals three features.

- The transfer is still direct, with no lunar flyby: the minimum spacecraft–Moon distance is 69 988 km, approximately 4 000 km outside the lunar sphere of influence and about 1 000 km closer than the Case 3a solution. The constraint pushes the trajectory marginally closer to the Moon but does not destabilise the no-flyby property.
- The two mid-course arrows (yellow and magenta nodes in Figure 5.9) visually overlap along the cislunar leg: this is the geometric counterpart of $\tau_2 = 0$, with the mid1 and mid2 nodes applied at the same point in space and time. The location of this compound node is near the outbound trajectory’s intersection with the lunar orbital plane, where the spacecraft velocity is sufficiently low to make the combined plane-and-energy manoeuvre energetically affordable.
- The segment durations are entirely rearranged with respect to Cases 1 and 3a: $\tau_1 = 3.76$ d (the longest of any case so far), $\tau_2 = 0$ d, $\tau_3 = 1.25$ d. The bulk of the

transfer is now the long backward coast on segment 1 from NRI to mid-course node, while the descent from that node to the LEO perigee occupies the remaining 1.25 d.

LEO → NRHO transfer (3 impulses) - Case 3b (Kourou $4^\circ < i < 8^\circ$) - Lunar-plane-aligned (rotated J2000)




Figure 5.9. Case 3b optimal transfer in the lunar orbital plane-aligned ECI frame, with the Kourou inclination constraint ($4^\circ \leq i \leq 8^\circ$) active.




Figure 5.10. Case 3b optimal transfer in the selenocentric frame.

5.3.3 ΔV Budget and the Compound Mid-course Manoeuvre

The total cost of the constrained solution is $J^* = 4.261$ km/s, a +478 m/s increase (+12.6%) with respect to the baseline. The decomposition is reported in Table 5.11 alongside the baseline and Cape Canaveral values.

Table 5.11. ΔV budget for Case 3b, compared with the baseline (Case 1) and with the Cape Canaveral solution (Case 3a).

Manoeuvre	Case 1 [m/s]	Case 3a [m/s]	Case 3b [m/s]	Δ (3b – 1) [m/s]
LEO insertion (TLI, Δv_{circ})	3041	3040	3013	–28
Mid-course correction 1	2	318	428	+426
Mid-course correction 2	0	1	347	+347
Compound mid (mid1 + mid2)	2	319	775	+773
NRHO insertion (Δv_{NRI})	740	470	473	–267
On-board total ΔV_{sc}	742	789	1248	+506
Total J^*	3783	3829	4261	+478

The redistribution exposes three distinct effects, in increasing order of operational importance.

The TLI cost decreases slightly. The trans-lunar injection cost $\Delta v_{\text{circ}} = 3013$ m/s is -28 m/s (-0.9%) below the baseline value. The effect is real but small, and originates in the slightly lower perigee velocity required by a trajectory whose initial plane is closer to the equator: a circular LEO at $i = 8^\circ$, when injected onto the transfer arc, yields a marginally lower asymptotic energy than a higher-inclination LEO. The contribution of this effect to the overall comparison with the baseline is negligible.

The NRHO insertion impulse does not scale with the misalignment angle. A naive scaling of the Cape Canaveral result by the plane-change ratio $\sin(\Delta i_{\text{Kou}}/2)/\sin(\Delta i_{\text{Cape}}/2) \approx 3.3$ would predict $\|\Delta \mathbf{v}_{\text{NRI}}\| \approx 1500$ m/s. The actual value of 473 m/s is essentially indistinguishable from the Cape Canaveral one (470 m/s). The NRHO insertion is therefore decoupled from the inclination penalty: the optimiser absorbs the entire additional plane-change effort in the mid-course nodes and arrives at the NRHO apolune with a trajectory geometry tailored to the same capture cost as in Case 3a. This is consistent with the design principle that out-of-plane corrections are performed where they are cheap (low velocity, mid-course), not where the spacecraft has already braked into the target orbit.

The compound mid-course manoeuvre carries the entire penalty. The dominant contribution to the Kourou penalty is the compound mid-course magnitude $\|\Delta \mathbf{v}_{\text{mid1}} + \Delta \mathbf{v}_{\text{mid2}}\| = 775$ m/s, $+773$ m/s above the baseline value and $+456$ m/s above the Cape Canaveral one. This compound vector, in the lunar orbital plane-aligned frame, is $[+0.394, +0.444, -0.498]$ km/s. Its magnitude exceeds the maximum admissible for a single impulse node under the per-component bound $|\Delta v_{\text{mid},k}| \leq 0.300$ km/s, which limits a single impulse to a magnitude of at most $\sqrt{3} \cdot 0.300 = 0.520$ km/s. The optimiser therefore splits the 775 m/s compound manoeuvre across the two available mid-course nodes, and co-locates them in time ($\tau_2 = 0$) so that the two impulses act on the same dynamical state and combine linearly into the single physically meaningful manoeuvre. The split between mid1 and mid2 is mathematically non-unique, as confirmed by the multi-start analysis of Section 5.3.4.

The decomposition of the four impulses in the local RTN frame is reported in Figure 5.11. Three observations follow.

- The TLI impulse remains essentially purely tangential, as in all previous cases.
- The NRI impulse retains the structure observed in Case 3a - predominantly tangential with a smaller normal contribution - consistent with its capture role and unchanged magnitude.
- The mid1 and mid2 bars must be read jointly. Their individual decompositions are arbitrary in the sense discussed above; the physically meaningful quantity is their sum. The compound mid-course manoeuvre is dominated by comparable contributions in the tangential and normal directions, with the normal component representing the bulk of the $\sim 33^\circ$ plane change and the tangential component compensating the energy mismatch of the constrained transfer plane.




Figure 5.11. Decomposition of the four impulses of Case 3b in the local Radial–Tangential–Normal (RTN) frame, computed at the pre-impulse spacecraft state.

A complementary view of the LEO end-state in the pure ECI J2000 frame is provided in Figure 5.12. The orbital inclination $i = 8.0^\circ$ saturates the upper bound of the Kourou range and lies $\sim 32.6^\circ$ below the unconstrained baseline value.




Figure 5.12. Case 3b LEO departure region in the ECI J2000 frame.

5.3.4 Multi-start Convergence, Active Constraints, and Operational Look-ahead

Multi-start convergence and the parametrisation degeneracy. The ten independent multi-start seeds yielded only three operationally-feasible solutions (Seeds 9, 7, 3 and 8); the remaining six violated the constraint set at the 10^{-3} tolerance, primarily through perigee-altitude or TOF infeasibilities. The summary of all ten outcomes is reported in Table 5.12; the lower convergence rate compared with Case 3a is consistent with the substantially tighter inclination requirement, which restricts the feasible region considerably.

The most informative outcome is that Seeds 9 and 7 converge to identical mission cost $J^* = 4.2615$ km/s and to identical $\Delta \mathbf{v}_{\text{NRI}}$, τ_1 and τ_{tot} , but to different individual splits of the compound mid-course vector between mid1 and mid2. The compound vector $\Delta \mathbf{v}_{\text{mid1}} + \Delta \mathbf{v}_{\text{mid2}}$ is, in both seeds, $[+0.394, +0.444, -0.498]$ km/s to four decimal places; the split is $[+0.218, +0.246, -0.276] + [+0.176, +0.199, -0.223]$ in Seed 9 and $[+0.157, +0.177, -0.198] + [+0.237, +0.267, -0.300]$ in Seed 7, with the second seed saturating its $-z$ component of mid2 against the lower bound. The two seeds are therefore two distinct realisations of the same physical optimum: this is direct empirical confirmation that the true decision variable, at convergence, is the compound mid vector rather than the individual impulses,

and that the three-impulse parametrisation is structurally redundant in this regime.

Table 5.12. Multi-start convergence of Case 3b (Kourou inclination constraint).

Seed	J [km/s]	τ_{tot} [d]	$\ \Delta \mathbf{v}_{\text{NRI}}\ $ [m/s]	compound mid [m/s]	h_p [km]	i [deg]	Feasibility
9	4.262	5.00	473	775	600.0	8.000	feas
7	4.262	5.00	473	775	600.0	8.000	feas (same as 9)
2	4.268	5.00	451	798	590.4	8.009	infeas
3	4.314	4.98	557	770	600.0	7.638	feas
10	4.433	5.56	743	664	600.0	8.006	infeas
4	4.546	8.09	724	890	619.1	7.790	infeas
5	4.728	10.37	937	825	581.8	8.159	infeas
6	4.743	14.69	835	907	600.0	8.253	infeas
8	5.035	5.00	1217	942	599.6	8.000	feas
1	5.072	5.00	1273	916	600.0	7.992	infeas

Active constraints at convergence. At the global optimum (Seed 9), all four active constraints of Case 3b are satisfied:

- **Perigee altitude:** $h_p = 600.00$ km, exactly at the upper bound of the admissible range.
- **Total time of flight:** $\tau_{\text{tot}} = 5.000$ d, exactly at the upper bound. As in Case 3a, the TOF constraint is marginally active.
- **LEO inclination:** $i = 8.000^\circ$, saturating the upper bound of the Kourou range $[4^\circ, 8^\circ]$. The optimiser pushes the inclination as high as the launch site allows in order to minimise the plane-change cost.
- **NRHO insertion:** matched to within $\|\mathbf{c}\|_\infty = 5.91 \cdot 10^{-11}$, well below both the operational and the strict feasibility tolerances.

Operational constraint look-ahead. The Van Allen radiation belt transit is $t_{\text{inner}} = 0.32$ h within the inner belt and $t_{\text{outer}} = 1.65$ h within the outer belt, both within their respective 0.50 h and 2.00 h operational thresholds that will be enforced in Case 4. The outer-belt transit is slightly shorter than in Cases 1 and 3a (1.73 and 1.75 h, respectively). The Case 4 radiation constraint set is therefore already satisfied by the Case 3b optimum, as for the Cape Canaveral case.

5.3.5 Comparison: Cape Canaveral vs Kourou

The Cape Canaveral and Kourou cases together quantify the impact of launch-site selection on a crewed LEO $\rightarrow$ NRHO transfer. The side-by-side comparison is presented in Table 5.13; three findings deserve a closing comment.

Table 5.13. Side-by-side comparison of the Cape Canaveral (Case 3a) and the Kourou (Case 3b) solutions, with the unconstrained baseline (Case 1) included for reference.

Quantity	Case 1 (baseline)	Case 3a (Cape)	Case 3b (Kourou)
LEO inclination i [deg]	40.6	31.0	8.0
Misalignment vs baseline Δi [deg]	0	9.6	32.6
TLI cost Δv_{circ} [m/s]	3041	3040	3013
Compound mid-course $\Sigma \Delta \mathbf{v}_{\text{mid}}$ [m/s]	2	319	775
NRHO insertion $\Delta \mathbf{v}_{\text{NRI}}$ [m/s]	740	470	473
On-board cost ΔV_{sc} [m/s]	742	789	1248
Total cost J^* [km/s]	3.783	3.829	4.261
On-board penalty vs baseline [m/s]	0	+47	+506
Total penalty vs baseline [m/s]	0	+46	+478

The launcher contribution is essentially insensitive to launch site. The trans-lunar injection cost differs by only -28 m/s between Kourou and Cape, with the equatorial Kourou solution slightly cheaper on TLI. The benefit associated with the equatorial launch site, which is decisive for geostationary-transfer missions, is here almost completely irrelevant: the $\sim 1\%$ saving on Δv_{circ} is dwarfed by the much larger on-board penalty discussed below.

The on-board penalty non-linear scaling with the misalignment angle. The misalignment ratio is $\Delta i_{\text{Kou}}/\Delta i_{\text{Cape}} = 32.6/9.6 \approx 3.4$, and the linearised plane-change formula would predict a ratio of compound mid-course magnitudes of $\sin(\Delta i_{\text{Kou}}/2)/\sin(\Delta i_{\text{Cape}}/2) \approx 3.3$. The observed ratio is $775/319 = 2.4$, in fact smaller than the linear prediction - a consequence of the optimiser having access to additional energy trade-offs through the tangential mid-course component. However, the on-board cost ratio $\Delta V_{\text{sc,Kou}}/\Delta V_{\text{sc,Cape}} = 1248/789 = 1.58$ understates the operational penalty, because what really matters for mission sizing is the incremental cost with respect to the unconstrained baseline. On that metric, the Kourou penalty $+506$ m/s is ~ 11 times the Cape Canaveral one $+47$ m/s: the plane-change ratio of 3.3 is amplified by the loss of the NRI-side compensation mechanism, which in Case 3a still allowed a -270 m/s recovery on the insertion impulse but in Case 3b is exhausted ($\Delta \mathbf{v}_{\text{NRI}}$ stays at the Case 3a value rather than dropping further). The Kourou geometry has, in this sense, exceeded the range over which the optimiser can compensate the additional plane-change demand by reshaping the rest of the trajectory.

Mission-design implication. For a crewed LEO $\rightarrow$ NRHO transfer in the framework of this work, Cape Canaveral is preferable to Kourou by an on-board ΔV margin of approximately 460 m/s, corresponding to a propellant mass fraction increase of order 15–20% for a typical hypergolic upper stage. The result is a direct consequence of the high inclination of the target NRHO with respect to the Earth equator, and inverts the launch-site ranking that holds for the geostationary-transfer scenario for which Kourou’s equatorial location is conventionally optimised. The two launch-site branches are nonetheless carried forward in parallel through the subsequent cases of this chapter (Van Allen belt transit in Case 4 and free NRHO arrival point in Case 5), so that the Kourou penalty can be assessed against each successive operational requirement and the conclusions of the present comparison can be revisited at the end of the cumulative-constraint sequence.

5.4 Effect of Van Allen Radiation Constraint

The fourth operational constraint of the cumulative sequence enforces upper bounds on the time spent by the spacecraft within the Earth’s inner and outer Van Allen radiation belts. The thresholds adopted in this work are derived directly from the flight measurements of the Artemis I mission [48]: $t_{\text{inner}} \leq 0.50$ h and $t_{\text{outer}} \leq 2.00$ h, corresponding to the cumulative transit times observed at the cabin-wall HSU2 location of the Orion vehicle (Extended Data Table 1 of the cited reference). The HSU2 sensor lies on the crew-cabin wall and represents the worst-case shielding for a typical crew workstation. The two thresholds are imposed separately on each belt rather than as a single cumulative bound, reflecting the qualitatively different radiation environments and biological risks of the two regions: the inner belt is dominated by high-energy protons that drive acute-dose risk, while the outer belt is composed predominantly of electrons whose contribution to the absorbed dose is mediated by secondary bremsstrahlung and is biologically less aggressive.

Two constraint sets are studied in this section. The operational set (Cases 4a and 4b) adopts the Artemis I thresholds directly. The tightened set (Cases 4c and 4d) further restricts the outer-belt transit to $t_{\text{outer}} \leq 1.50$ h, a 25% reduction with respect to the Artemis I value; the inner-belt threshold is retained at 0.50 h. The motivation for tightening the outer-belt threshold, rather than the more biologically aggressive inner-belt one, is methodological: the outer-belt transit at the Case 4a/4b optima ($t_{\text{outer}} \approx 1.65\text{--}1.68$ h, see Sections 5.4.1 and 5.4.2) is non-binding with a margin of approximately 17–19% relative to the operational threshold of 2.00 h. Reducing the threshold to 1.50 h transitions this constraint from non-binding to active, allowing the direct estimation of the ΔV cost incurred when the outer-belt transit becomes a hard limit on the optimisation. Both constraint sets are applied to both launch-site branches (suffix “a” and “c” for Cape Canaveral, “b” and “d” for Kourou), giving four sub-cases analysed in Sections 5.4.1–5.4.4, with a comparative discussion in Section 5.10.

5.4.1 Operational Thresholds: Cape Canaveral (Case 4a)

The Case 4a optimisation activates the operational Van Allen thresholds on top of the Case 3a constraint set. The best operationally-feasible seed converges to $J^* = 3.836$ km/s at $\tau_{\text{tot}} = 4.79$ d, $h_p = 600$ km and $i = 31.0^\circ$, with the radiation transit at $t_{\text{inner}} = 0.32$ h (64% of the threshold) and $t_{\text{outer}} = 1.65$ h (83% of the threshold), as shown in Figure 5.13. Both thresholds are satisfied with substantial margin, and the radiation constraint set is therefore operationally non-binding for the Cape Canaveral global optimum.




Figure 5.13. Spacecraft altitude versus time of flight for the Case 4a optimum (Seed 6).

A direct comparison with the Case 3a outcome is reported in Table 5.14. The Case 4a multi-start converges to a solution at +7 m/s (+0.18%) above the Case 3a value, with a substantial redistribution of the on-board budget between the two mid-course nodes: $\|\Delta \mathbf{v}_{\text{mid1}}\|$ collapses from 318 m/s in Case 3a to 14 m/s in Case 4a, while $\|\Delta \mathbf{v}_{\text{mid2}}\|$ grows from 1 m/s to 243 m/s. The compound mid-course magnitude $\|\Delta \mathbf{v}_{\text{mid1}} + \Delta \mathbf{v}_{\text{mid2}}\|$ remains comparable in the two cases (318 m/s in Case 3a versus 256 m/s in Case 4a), and the position of the dominant mid-course node along the trajectory is also similar (mid-coast in both cases, applied at a comparable spacecraft state). The redistribution is therefore best interpreted as an alternative parametric labelling of essentially the same physical plane-change manoeuvre, analogous to the mid1/mid2 split degeneracy already documented for Case 3b in Section 5.3.4, rather than as a fundamentally different trajectory architecture.

Table 5.14. Comparison of the Case 3a and Case 4a optima.

Quantity	Case 3a	Case 4a	Δ
J^* [km/s]	3.829	3.836	+0.007
τ_{tot} [d]	5.00	4.79	-0.21
$\ \Delta \mathbf{v}_{\text{NRI}}\ $ [m/s]	470	539	+69
$\ \Delta \mathbf{v}_{\text{mid1}}\ $ [m/s]	318	14	-304
$\ \Delta \mathbf{v}_{\text{mid2}}\ $ [m/s]	1	243	+242
Compound mid $\ \Delta \mathbf{v}_{\text{mid1}} + \Delta \mathbf{v}_{\text{mid2}}\ $ [m/s]	318	256	-62
ΔV_{sc} [m/s]	789	796	+7
t_{inner} [h]	0.32	0.32	0
t_{outer} [h]	1.75	1.65	-0.10

A subtle but documented side effect of activating the radiation constraint, even though non-binding at the optimum, is a 6% reduction of the outer-belt transit time relative to Case 3a (1.75 $\rightarrow$ 1.65 h). The penalty-based formulation of the constraint, intrinsic to the PSO global search, appears to bias the multi-start exploration towards solutions

with larger constraint margins even when the constraint itself does not formally activate. The effect is operationally inconsequential – +7 m/s of cost in exchange for 0.10 h of additional outer-belt margin – but illustrates that the cumulative-constraint sequence affects the optimisation landscape in non-trivial ways beyond the simple binding/non-binding dichotomy.

For the cumulative-constraint analysis of this chapter, the conclusion is that the operational Van Allen constraint set adds no measurable ΔV penalty to the Cape Canaveral global optimum beyond what is already captured by Case 3a, in close analogy to the role of the time-of-flight constraint in Case 2.

5.4.2 Operational Thresholds: Kourou (Case 4b)

The Case 4b optimisation activates the operational Van Allen thresholds on top of the Case 3b constraint set, with the launch-site inclination restricted to the Kourou window $4^\circ \leq i \leq 8^\circ$. The convergence of the multi-start is substantially weaker than in Case 4a: only one seed out of ten satisfies the operational feasibility tolerance $\|\mathbf{c}\|_\infty \leq 10^{-3}$ (Seed 4, with $\|\mathbf{c}\|_\infty = 5.12 \cdot 10^{-4}$). This reduced convergence rate is consistent with the much tighter feasible region of the Kourou case already documented in Section 5.3.4, where the inclination constraint alone restricted the multi-start to four feasible solutions out of ten; the addition of the radiation constraint further reduces the feasible domain accessible to the solver.

The best operationally-feasible seed (Seed 4) converges to $J^* = 4.266$ km/s at $\tau_{\text{tot}} = 5.00$ d, $h_p = 600$ km and $i = 8.00^\circ$, with the radiation transit at $t_{\text{inner}} = 0.31$ h (62% of the threshold) and $t_{\text{outer}} = 1.68$ h (84% of the threshold), as shown in Figure 5.14. Both thresholds are satisfied with margin, and the radiation constraint set is therefore operationally non-binding also for the Kourou global optimum.




Figure 5.14. Spacecraft altitude versus time of flight for the Case 4b optimum (Seed 4).

A direct comparison with the Case 3b outcome is reported in Table 5.15. The Case 4b

multi-start converges to a solution at +4 m/s (+0.09%) above the Case 3b value, with a near-identical compound mid-course magnitude ($775 \rightarrow 778$ m/s) split slightly differently between mid1 and mid2. The most informative difference is the segment-2 duration: in Case 3b the two mid-course nodes were exactly co-located in time ($\tau_2 = 0$ d), whereas in Case 4b they are separated by $\tau_2 = 0.44$ d. The physical compound manoeuvre remains identical to four significant figures ($\|\Delta\mathbf{v}_{\text{mid1}} + \Delta\mathbf{v}_{\text{mid2}}\| = 0.78$ km/s in both cases), and the NRHO insertion impulse is unchanged at $\|\Delta\mathbf{v}_{\text{NRI}}\| = 474$ m/s. This is a third independent demonstration that the true decision variable, at convergence, is the compound mid-course vector rather than the individual impulses.

Table 5.15. Comparison of the Case 3b and Case 4b optima.

Quantity	Case 3b	Case 4b	Δ
J^* [km/s]	4.262	4.266	+0.004
τ_{tot} [d]	5.00	5.00	0
τ_2 [d]	0.000	0.439	+0.44
$\ \Delta\mathbf{v}_{\text{NRI}}\ $ [m/s]	473	474	+1
$\ \Delta\mathbf{v}_{\text{mid1}}\ $ [m/s]	428	407	-21
$\ \Delta\mathbf{v}_{\text{mid2}}\ $ [m/s]	347	373	+26
Compound mid $\ \Delta\mathbf{v}_{\text{mid1}} + \Delta\mathbf{v}_{\text{mid2}}\ $ [m/s]	775	778	+3
ΔV_{sc} [m/s]	1248	1254	+6
t_{inner} [h]	0.32	0.31	-0.01
t_{outer} [h]	1.65	1.68	+0.03

The Van Allen transit times of the Case 4b optimum are essentially indistinguishable from those of the Case 3b solution (± 0.03 h on each belt), in contrast with the small but systematic outer-belt reduction documented for Case 4a relative to Case 3a (Section 5.4.1). The dispersion of the two cases is therefore fully attributable to the stochasticity of the PSO multi-start rather than to a constraint-induced bias on the search. This is consistent with the much tighter feasible region of the Kourou problem, in which the optimiser has substantially less room to manoeuvre once the inclination constraint is already binding from above and the compound mid-course vector is essentially fixed by the plane-change requirement.

For the cumulative-constraint analysis of this chapter, the conclusion is that the operational Van Allen constraint set adds no measurable ΔV penalty to the Kourou global optimum beyond what is already captured by Case 3b, in close analogy to the Cape Canaveral case. The non-binding character of the thresholds derived from Artemis I mission is therefore confirmed for both launch-site branches. The tightened outer-belt analysis of Cases 4c and 4d, reported in the following sections, examines the additional ΔV cost that would be incurred if the outer-belt threshold were reduced from 2 h to 1.5 h.

5.4.3 Tightened Outer Belt: Cape Canaveral (Case 4c)

The Case 4c optimisation activates the tightened Van Allen outer-belt threshold $t_{\text{outer}} \leq 1.50$ h on top of the Case 3a constraint set; the inner-belt threshold is retained at 0.50 h, consistently with Case 4a. The reduction from 2.00 h to 1.50 h transitions the outer-belt limit from non-binding to active: the Case 4a optimum, at $t_{\text{outer}} = 1.65$ h, is no longer

feasible under the tightened threshold, and the multi-start must converge to a trajectory family with a shorter outer-belt transit.

Selection criterion. Under the tightened threshold no seed of the multi-start satisfies the operational feasibility tolerance $\|\mathbf{c}\|_\infty \leq 10^{-3}$ adopted for Cases 4a/4b. The reported solution is therefore selected as the seed with the lowest J^* whose dominant active constraint at convergence is the tightened outer-belt threshold (rather than a different constraint of the cumulative set). This criterion ensures that the solution captures the direct ΔV cost of the tightening, rather than a parametric escape into a different active-constraint regime; the residual constraint excess is reported explicitly as part of the solution metadata.

The selected solution converges to $J^* = 3.866$ km/s at $\tau_{\text{tot}} = 4.79$ d, $h_p = 530.9$ km and $i = 30.44^\circ$, with the outer-belt transit at $t_{\text{outer}} = 1.61$ h. The residual constraint excess is $t_{\text{outer}} - 1.50 = 0.113$ h = 6.78 min, corresponding to an operational feasibility violation of $\|\mathbf{c}\|_\infty = 1.13 \cdot 10^{-1}$. The inner-belt transit remains well below the 0.50 h operational threshold. The altitude profile of the optimum is shown in Figure 5.15.




Figure 5.15. Spacecraft altitude versus time of flight for the Case 4c optimum.

A direct comparison with the Case 4a optimum is reported in Table 5.16. The tightened outer-belt threshold drives a +30 m/s increase in J^* (+0.78%), concentrated almost entirely in the NRHO-insertion impulse magnitude, which grows from 539 m/s to 626 m/s (+87 m/s). The compound mid-course magnitude is reduced by -72 m/s ($256 \rightarrow 184$ m/s), and the time-of-flight is essentially unchanged at 4.79 d. The geometry of the transfer is structurally preserved: the trajectory family is the same as in Case 4a, but with an additional energetic effort applied at the NRHO terminal node to enforce a more rapid traversal of the outer-belt region.

Table 5.16. Comparison of the Case 4a (operational) and Case 4c (tightened outer-belt) optima for the Cape Canaveral branch.

Quantity	Case 4a	Case 4c	Δ
J^* [km/s]	3.836	3.866	+0.030
τ_{tot} [d]	4.79	4.79	0
h_p [km]	600	531	−69
i [deg]	31.0	30.4	−0.6
$\ \Delta \mathbf{v}_{\text{NRI}}\ $ [m/s]	539	626	+87
$\ \Delta \mathbf{v}_{\text{mid1}}\ $ [m/s]	14	120	+106
$\ \Delta \mathbf{v}_{\text{mid2}}\ $ [m/s]	243	64	−179
Compound mid $\ \Delta \mathbf{v}_{\text{mid1}} + \Delta \mathbf{v}_{\text{mid2}}\ $ [m/s]	256	184	−72
ΔV_{sc} [m/s]	796	810	+14
t_{outer} [h]	1.65	1.61	−0.04
$\ \mathbf{c}\ _{\infty}$	$\leq 10^{-3}$	$1.13 \cdot 10^{-1}$	—

The residual constraint excess of 0.113 h, which exceeds the operational feasibility tolerance by approximately two orders of magnitude, indicates that the 1.50 h outer-belt limit lies close to the boundary of the feasible region accessible to the PSO+SQP multi-start with the standard 10-seed compute budget. The result is therefore best interpreted as a lower bound on the actual ΔV cost of the tightening, since a more thoroughly sampled multi-start might yield strictly feasible solutions at slightly higher J^* . The methodological implications of this boundary behaviour, together with the analogous result for the Kourou branch (Case 4d), are discussed in Section 5.10.

5.4.4 Tightened Outer Belt: Kourou (Case 4d)

The Case 4d optimisation activates the tightened Van Allen outer-belt threshold $t_{\text{outer}} \leq 1.50$ h on top of the Case 3b constraint set, with the launch-site inclination restricted to the Kourou window $4^\circ \leq i \leq 8^\circ$. The convergence of the multi-start is qualitatively weaker than in Case 4c: as anticipated by the reduced feasibility documented for the Kourou branch under the operational Van Allen thresholds (Section 5.4.2), the further tightening on t_{outer} substantially restricts the feasible region accessible to the solver.

Selection criterion. Under the tightened threshold no seed of the multi-start satisfies the operational feasibility tolerance $\|\mathbf{c}\|_{\infty} \leq 10^{-3}$ adopted for Cases 4a/4b. As in Case 4c, the reported solution is selected as the seed with the lowest J^* whose dominant active constraint at convergence is the tightened outer-belt threshold. For the Kourou branch this filter is more selective than for the Cape Canaveral branch: several seeds of the multi-start converge to trajectories that satisfy the outer-belt constraint marginally but violate a different cumulative constraint instead (typically the time-of-flight $\tau_{\text{tot}} \leq 5$ d, which becomes active due to the geometric coupling between the narrow Kourou inclination window and the tightened outer-belt transit time). Such convergence modes do not represent the cost of the outer-belt tightening per se and are therefore excluded from the reported solution; their existence is discussed quantitatively in Section 5.10.

The selected solution converges to $J^* = 4.394$ km/s at $\tau_{\text{tot}} = 5.00$ d, $h_p = 324.0$ km and $i = 7.94^\circ$, with the outer-belt transit at $t_{\text{outer}} = 1.62$ h. The residual constraint excess is $t_{\text{outer}} - 1.50 = 0.125$ h = 7.49 min, corresponding to a feasibility violation of

$\|\mathbf{c}\|_\infty = 1.25 \cdot 10^{-1}$. The time-of-flight constraint is active ($\tau_{\text{tot}} = 4.999$ d, 0.04 min below the 5 d limit), and the inclination is essentially at the upper Kourou bound ($i = 7.94^\circ$ vs $i_{\text{max}} = 8^\circ$). The altitude profile of the optimum is shown in Figure 5.16.




Figure 5.16. Spacecraft altitude versus time of flight for the Case 4d optimum.

A direct comparison with the Case 4b optimum is reported in Table 5.17. The tightened outer-belt threshold drives a +128 m/s increase in J^* (+3.00%), substantially larger than the +30 m/s observed for the Cape Canaveral branch (Case 4c). The cost is again concentrated at the NRHO-insertion impulse, which grows from 474 m/s to 572 m/s (+98 m/s), while the mid-course budget is essentially preserved. The time-of-flight transitions from non-binding to active, and the perigee altitude collapses from $h_p = 600$ km to $h_p = 324$ km, indicating a substantially different geometry of the descent towards LEO.

Table 5.17. Comparison of the Case 4b (operational) and Case 4d (tightened outer-belt) optima for the Kourou branch.

Quantity	Case 4b	Case 4d	Δ
J^* [km/s]	4.266	4.394	+0.128
τ_{tot} [d]	5.00	5.00	0
h_p [km]	600	324	-276
i [deg]	8.00	7.94	-0.06
$\ \Delta\mathbf{v}_{\text{NRI}}\ $ [m/s]	474	572	+98
$\ \Delta\mathbf{v}_{\text{mid1}}\ $ [m/s]	407	396	-11
$\ \Delta\mathbf{v}_{\text{mid2}}\ $ [m/s]	373	345	-28
Compound mid $\ \Delta\mathbf{v}_{\text{mid1}} + \Delta\mathbf{v}_{\text{mid2}}\ $ [m/s]	778	741	-37
ΔV_{sc} [m/s]	1254	1313	+59
t_{outer} [h]	1.68	1.62	-0.06
$\ \mathbf{c}\ _\infty$	$\leq 10^{-3}$	$1.25 \cdot 10^{-1}$	—

The +128 m/s cost of the outer-belt tightening for the Kourou branch is qualitatively

larger than the +30 m/s observed for the Cape Canaveral branch because the additional constraint interacts with the already-binding inclination constraint of the Kourou window. The combined geometric requirement of (i) a plane change from $i \in [4^\circ, 8^\circ]$ to the NRHO inclination, (ii) a time-of-flight under 5 d, and (iii) an outer-belt transit under 1.50 h leaves the optimiser with very limited room to manoeuvre, and the converged solution sits at the simultaneous activation of all three constraints. As for Case 4c, the residual constraint excess of 0.125 h indicates that the reported cost is best interpreted as a lower bound on the actual cost of the tightening, with the full methodological discussion deferred to Section 5.10.

5.5 Effect of the Free NRHO Arrival Point: Cape Canaveral (Case 5a)

5.5.1 Motivation and parametrisation

The cumulative-constraint sequence of Sections 5.1–5.4.1 has, throughout, kept the NRHO arrival phase fixed at $\alpha_{\text{apo}} = 0.8539$, i.e. exactly at the apolune passage of the reference periodic orbit. This choice mirrors the early literature on direct LEO $\rightarrow$ NRHO transfers, in which apolune insertion is preferred on navigational grounds: the spacecraft–Moon distance is largest, the local NRHO velocity is smallest, and the sensitivity of the orbital state to insertion errors is correspondingly minimised [13]. The apolune lock, however, removes one degree of freedom from the geometric matching problem, and the optimiser has been forced to absorb the entire mismatch between the trans-lunar arc geometry and the target NRHO state into the capture impulse $\Delta \mathbf{v}_{\text{NRI}}$ and the mid-course correction architecture.

Case 5a relaxes this restriction. The NRHO arrival phase $\alpha \in [0, 1)$, expressed in the perilune-to-perilune coordinate of the reference orbit, is added to the decision vector as a free variable; the reference NRHO state is recomputed at the optimised α along the ephemeris-propagated orbit. The physical interpretation is the following: the optimiser is now allowed to choose the insertion point along the NRHO that minimises the total mission cost, subject to all four operational constraints introduced in Cases 2–4a ($\tau_{\text{tot}} \leq 5$ d, $h_p \in [150 \text{ km}, 600 \text{ km}]$, $i \in [26^\circ, 31^\circ]$ for Cape Canaveral, and the Artemis I–derived Van Allen thresholds $t_{\text{inner}} \leq 0.50$ h and $t_{\text{outer}} \leq 2.00$ h).

The decision vector has dimension 13 rather than 12:

$$\mathbf{x} = [\Delta \mathbf{v}_{\text{NRI}}; \tau_1; \Delta \mathbf{v}_{\text{mid1}}; \tau_2; \Delta \mathbf{v}_{\text{mid2}}; \tau_{3,\text{max}}; \alpha] \in \mathbb{R}^{13}, \quad (5.2)$$

where $\tau_{3,\text{max}}$ is the upper bound on the duration of segment 3 (which is event-terminated by the perigee-crossing detector). On the basis of the navigational argument recalled above, the expectation is that the optimiser will settle on a solution close to apolune ($\alpha \approx \alpha_{\text{apo}}$), with at most a modest displacement to exploit favourable geometric alignment with the trans-lunar arc.

A note on numerical tolerances. The operational feasibility tolerance for Case 5a is relaxed to $\|\mathbf{c}\|_\infty \leq 2 \times 10^{-2}$, compared with 10^{-3} in Cases 3a–4a. The relaxation is a

Table 5.18. Optimal decision vector $\mathbf{x}^*$ for Case 5a (Cape Canaveral, $26^\circ \leq i \leq 31^\circ$, Van Allen radiation constraint, free NRHO arrival point).

Variable	Description	Value	Unit
$\Delta v_{\text{NRI},x}$	NRHO insertion impulse, x	+0.0056	km/s
$\Delta v_{\text{NRI},y}$	NRHO insertion impulse, y	+0.3419	km/s
$\Delta v_{\text{NRI},z}$	NRHO insertion impulse, z	−0.0321	km/s
τ_1	Duration, segment 1	0.738	d
$\Delta v_{\text{mid1},x}$	Mid-course 1, x	+0.0157	km/s
$\Delta v_{\text{mid1},y}$	Mid-course 1, y	+0.3000	km/s
$\Delta v_{\text{mid1},z}$	Mid-course 1, z	−0.0674	km/s
τ_2	Duration, segment 2	0.223	d
$\Delta v_{\text{mid2},x}$	Mid-course 2, x	−0.0067	km/s
$\Delta v_{\text{mid2},y}$	Mid-course 2, y	+0.1255	km/s
$\Delta v_{\text{mid2},z}$	Mid-course 2, z	−0.0332	km/s
$\tau_{3,\text{max}}$	Upper bound, segment 3	8.732	d
α	NRHO arrival phase (free)	0.4376	–
<i>Derived quantities and diagnostics</i>			
$\ \Delta \mathbf{v}_{\text{NRI}}\ $	NRHO insertion magnitude	0.3434	km/s
$\ \Delta \mathbf{v}_{\text{mid1}}\ $	Mid-course 1 magnitude	0.3079	km/s
$\ \Delta \mathbf{v}_{\text{mid2}}\ $	Mid-course 2 magnitude	0.1300	km/s
$\ \Delta \mathbf{v}_{\text{mid1}} + \Delta \mathbf{v}_{\text{mid2}}\ $	Compound mid magnitude	0.438	km/s
J^*	Total mission cost	3.821	km/s
τ_{tot}	Total realised flight time	≈ 4.99	d
$\Delta \alpha = \alpha_{\text{apo}} - \alpha^*$	Displacement from apolune	0.0596	–
Δt_{apo}	Equivalent time before apolune	9.4	h
h_p	LEO perigee altitude	578	km
i	LEO inclination	31.1	deg
$d_{\text{Moon}, \text{min}}$	Minimum spacecraft–Moon distance	69 632	km
$\ \mathbf{c}\ _\infty$	Maximum constraint violation	3.9×10^{-5}	–

precaution against the additional dimension of the decision space and the consequent risk of multi-start outliers landing in marginally infeasible regions of the cost landscape.

5.5.2 Optimal decision vector

The hybrid PSO–SQP scheme was executed over twenty independent multi-start seeds. The best operationally-feasible seed (Seed 18) converges to the solution summarised in Table 5.18, with a total mission cost $J^* = 3.821 \text{ km s}^{-1}$, a total realised flight time $\tau_{\text{tot}} = \tau_1 + \tau_2 + \tau_3 \approx 4.99 \text{ d}$ (segment 3 event-terminated by the perigee-crossing detector at $\tau_3 < \tau_{3,\text{max}}$), and an NRHO arrival phase $\alpha^* = 0.4376$.

The proximity of the optimal arrival point to the apolune is the central qualitative finding of this section, and is examined in detail in Section 5.5.4. The interpretation of $\alpha^* = 0.4376$ is the following: in the perilune-to-perilune normalisation of the NRHO, the apolune passage occurs at $\alpha_{\text{apo}} = 0.4971$; the optimal insertion is therefore $\Delta \alpha = 0.0596$ short of the apolune. Translated to absolute time along the NRHO motion (period $T_{\text{NRHO}} \approx 6.531 \text{ d}$), the insertion is performed approximately 9.4 h ahead of the apolune

passage: after the capture impulse, the spacecraft coasts naturally along the NRHO and reaches the apolune ≈ 9.4 h later.

The capture impulse magnitude $\|\Delta\mathbf{v}_{\text{NRI}}\| = 343 \text{ m s}^{-1}$ is markedly smaller than the corresponding Case 4a value of 539 m s^{-1} , and the compound mid-course magnitude $\|\Delta\mathbf{v}_{\text{mid1}} + \Delta\mathbf{v}_{\text{mid2}}\| = 438 \text{ m s}^{-1}$ is correspondingly larger than the Case 4a value of 256 m s^{-1} . The TLI cost $\Delta v_{\text{circ}} = 3039 \text{ m s}^{-1}$ is essentially indistinguishable from the launch-site-constrained values of the previous sections (3040 m s^{-1} in Case 3a, 3041 m s^{-1} in Case 4a). The redistribution between the NRI and the mid-course nodes deserves a careful comment, because, as documented in Section 5.6.1, Seed 18 is one realisation of a family of solutions that present the same cost but differ in their individual splitting of the on-board ΔV between the three impulse nodes. The physically meaningful quantity is the total on-board budget $\Delta V_{\text{sc}} = 781 \text{ m s}^{-1}$, not its per-node decomposition.

5.5.3 Trajectory geometry

The optimised trajectory is displayed in the lunar orbital plane-aligned ECI frame in Figure 5.17. The qualitative comparison with the Case 4a apolune-locked solution reveals two features.

- The insertion point lies on the inbound leg of the NRHO, short of the apolune passage. In Figure 5.17, the marker corresponding to the capture impulse falls before, rather than at, the maximum spacecraft–Moon distance along the displayed NRHO half-period. The free arrival point acts as a refinement of the geometric matching with the trans-lunar arc, not as a relocation to a fundamentally different region of the orbit.
- The transfer retains its direct, no-flyby character. The minimum spacecraft–Moon distance is $d_{\text{Moon, min}} = 69\,632 \text{ km}$, $\approx 3600 \text{ km}$ outside the lunar sphere of influence ($66\,000 \text{ km}$), and comparable to the Case 4a value (Section 5.4.1). No qualitatively new geometry such as a powered lunar flyby or a manifold-assisted insertion emerges from releasing the arrival phase.




Figure 5.17. Case 5a optimal transfer in the lunar orbital plane-aligned ECI frame, with the free NRHO arrival point active.

The Van Allen radiation constraint is satisfied with margin: $t_{\text{inner}} = 0.32$ h (64% of the 0.50 h threshold) and $t_{\text{outer}} = 1.79$ h (89% of the 2.00 h threshold). The outer-belt transit is slightly longer than in Case 4a (+0.14 h), reflecting the marginally different LEO end-state of the free-arrival-point solution (574 km instead of 600 km); both thresholds remain operationally non-binding.

5.5.4 ΔV budget and the marginal benefit of the free arrival point

The total cost $J^* = 3.821 \text{ km s}^{-1}$ lies -15 m s^{-1} (-0.4%) below the Case 4a value (3.836 km s^{-1}) and -8 m s^{-1} (-0.2%) below the Case 3a value (3.829 km s^{-1}). The detailed decomposition is reported in Table 5.19.

Table 5.19. ΔV budget for Case 5a, compared with the Cape Canaveral baseline of Case 3a (apolune insertion, no Van Allen constraint) and with Case 4a (apolune insertion, Van Allen constraint active).

Manoeuvre	Case 3a [m/s]	Case 4a [m/s]	Case 5a [m/s]	Δ (5a-4a) [m/s]
LEO insertion (TLI, Δv_{circ})	3 040	3 041	3 039	-2
Mid-course correction 1	318	14	308	+294
Mid-course correction 2	1	243	130	-113
Compound mid $\ \Delta \mathbf{v}_{\text{mid1}} + \Delta \mathbf{v}_{\text{mid2}}\ $	319	256	438	+182
NRHO insertion (Δv_{NRI})	470	539	343	-196
On-board total ΔV_{sc}	789	796	781	-15
Total J^*	3 829	3 836	3 821	-15

Three findings deserve comment.

The marginal benefit of the free arrival point is small in absolute terms. The on-board ΔV saving with respect to Case 4a is -15 m s^{-1} (-1.9%). Compared with the 47 m s^{-1} penalty associated with the Cape Canaveral inclination constraint of Case 3a, or with the 506 m s^{-1} penalty incurred in the Kourou case (Section 5.3), the recovery achieved by releasing the arrival phase is operationally minor. This is consistent with the geometric intuition that the apolune is a near-optimal insertion point: the local sensitivity of the cost to the arrival phase is small in a neighbourhood of α_{apo} , and the optimiser exploits only a fraction of this neighbourhood.

The optimum settles close to, but not exactly at, apolune. The displacement $\Delta \alpha = 0.0596$, equivalent to $\approx 9.4 \text{ h}$ or 6% of one full NRHO period, places the insertion point on the inbound leg of the orbit: after the capture impulse, the spacecraft coasts naturally for $\approx 9.4 \text{ h}$ before reaching the orbit's maximum lunar distance. The displacement is small enough that the navigational arguments in favour of apolune insertion (low local velocity, low sensitivity) remain qualitatively applicable; yet the value is consistent across the multi-start seeds (see Section 5.6.1) and statistically distinguishable from zero, with a multi-start standard deviation of only $\sigma_\alpha = 0.0035$ ($\approx 0.5 \text{ h}$) over the top five seeds. The fact that the optimiser settles short of the apolune - i.e. on the inbound leg rather than on the outbound one - is a non-trivial diagnostic of the symmetry-breaking introduced by the ephemeris model: in a perfectly periodic CR3BP analog of the NRHO, the cost landscape would be symmetric about apolune, and no preferred direction of displacement would emerge.

The redistribution of the on-board budget. The capture impulse $\|\Delta \mathbf{v}_{\text{NRI}}\| = 343 \text{ m s}^{-1}$ is -196 m s^{-1} smaller than in Case 4a, and the saving is more than offset by a 182 m s^{-1} growth of the compound mid-course magnitude. The total on-board budget is only -15 m s^{-1} lower than in Case 4a. The mid-course nodes therefore absorb additional ΔV that the apolune-locked solution had concentrated at the NRI. Crucially, as documented in Section 5.6.1, the global optimum cluster contains several distinct partitions of essentially the same total cost between the NRI and the mid-course nodes: Seed 18 ($\|\Delta \mathbf{v}_{\text{NRI}}\| = 343 \text{ m s}^{-1}$, compound mid = 438 m s^{-1}) and Seeds 19/17 ($\|\Delta \mathbf{v}_{\text{NRI}}\| \approx 469 \text{ m s}^{-1}$, compound mid $\approx 314 \text{ m s}^{-1}$, with mid2 reduced to a negligible 0 m s^{-1} to 2 m s^{-1}) converge to J^* within 2 m s^{-1} of one another but to clearly different impulse architectures.

This is structurally analogous to the mid1/mid2 split degeneracy already documented for Case 3b (Section 5.3.3), and it implies that the per-node decomposition reported in Table 5.18 should be interpreted as one of several realisations of the same physical optimum that present the same overall cost, rather than as a uniquely determined manoeuvre profile.

The decomposition of the four impulses in the local Radial-Tangential-Normal (RTN) frame is reported in Figure 5.18, computed at the pre-impulse spacecraft state. The TLI impulse is essentially purely tangential ($T \approx 3044 \text{ m s}^{-1}$, $R, N \approx 0$), as expected for an optimal energy-injection burn at perigee. The NRI impulse is dominated by the tangential component ($T \approx 320 \text{ m s}^{-1}$), consistent with its primary role as a braking manoeuvre along the NRHO tangent direction at the arrival point, with smaller radial and normal contributions that account for the residual mismatch between the trans-lunar arc and the locally aligned NRHO frame. The mid-course nodes mid1 and mid2 are predominantly tangential as well, with smaller normal components: their combined effect along the cislunar leg corresponds to a coupled energy and plane-correction manoeuvre, the same dynamical role observed in Case 3a (Section 5.2), but with the redistribution made possible by the free arrival point shifting part of the work from the NRI to the cislunar coast.




Figure 5.18. Decomposition of the four impulses of Case 5a in the local Radial-Tangential-Normal (RTN) frame, computed at the pre-impulse spacecraft state.

5.6 Sensitivity Analysis




Figure 5.19. Sensitivity analysis for Case 5a

The sensitivity of the Case 5a optimum to the width of the `idx_norm` bound is examined in Figure 5.19 by progressively relaxing the admissible offset from 0 (recovering the apolune-fixed Case 4a, $J^* = 3.835$ km/s) up to ± 0.5 . The converged ΔV decreases monotonically and saturates once the bound exceeds $|\text{idx}_{\text{opt}}| = 0.0662$, confirming that the free-apolune optimum of Case 5a ($J^* = 3.819$ km/s) is the true unconstrained minimum and not an artefact of the specific bound adopted.

5.6.1 Multi-start convergence, active constraints, and operational implications

Multi-start convergence. The twenty independent multi-start seeds yielded seventeen operationally-feasible solutions ($\|\mathbf{c}\|_{\infty} \leq 2 \times 10^{-2}$), of which sixteen satisfy the much stricter $\|\mathbf{c}\|_{\infty} \leq 10^{-3}$ tolerance, demonstrating that the relaxation of the operational tolerance for Case 5a is precautionary rather than necessary. The three infeasible seeds (Seeds 10, 16, 5) reach significantly higher maximum violations (in the range 0.2 to 8.0, dominated by perigee-altitude or TOF infeasibilities) and converge to costs $J > 4.12 \text{ km s}^{-1}$, well outside the feasible cluster. The summary of the ten lowest-cost outcomes is reported in Table 5.20.

Table 5.20. Multi-start convergence of Case 5a (Cape Canaveral, Van Allen, free NRHO arrival point), reporting the ten lowest-cost seeds.

Seed	J [km/s]	$\ \mathbf{c}\ _\infty$	α	$\ \Delta v_{\text{NRI}}\ $ [m/s]	$\ \Delta v_{\text{mid1}}\ $ [m/s]	$\ \Delta v_{\text{mid2}}\ $ [m/s]	comp. mid [m/s]
18	3.8207	$+3.9 \times 10^{-5}$	0.4376	343	308	130	438
19	3.8224	$+1.3 \times 10^{-6}$	0.4365	470	313	0	313
17	3.8224	$+1.1 \times 10^{-6}$	0.4372	468	313	2	315
13	3.8228	$+2.4 \times 10^{-4}$	0.4414	381	307	95	401
2	3.8272	$+2.1 \times 10^{-5}$	0.4306	601	143	44	185
3	3.8343	$+5.2 \times 10^{-6}$	0.4013	741	23	31	53
4	3.8480	-1.8×10^{-4}	0.4634	648	85	78	153
20	3.8629	$+1.3 \times 10^{-4}$	0.6320	520	161	143	303
15	3.8702	$+3.2 \times 10^{-4}$	0.4046	776	50	7	56
12	3.8741	$+1.4 \times 10^{-4}$	0.4302	561	251	21	270

Two diagnostics emerge from the table.

- The convergence cluster of the arrival phase is robust. The five lowest-cost seeds (Seeds 18, 19, 17, 13, 2) converge to $\alpha \in [0.4306, 0.4414]$ with a mean $\bar{\alpha} = 0.4367$ and a standard deviation $\sigma_\alpha = 0.0035$, equivalent to a temporal scatter of only ± 0.5 h around the mean displacement of 9.4 h before apolune. The displacement is therefore a robust property of the global optimum, not an artefact of the specific seed selected. Seed 3 ($\alpha = 0.4013$) lies marginally outside this cluster and converges to a slightly higher cost ($+13 \text{ m s}^{-1}$), consistent with the cost increasing as α moves further from $\bar{\alpha}$.
- The arrival-phase landscape contains a secondary basin far from apolune. Seeds 14 ($\alpha = 0.908$), 1 ($\alpha = 0.599$), 9 ($\alpha = 0.522$) and 11 ($\alpha = 0.173$) converge to feasible solutions with α in the range 0.17 to 0.91 and costs $J \in [3.89 \text{ km s}^{-1}, 4.09 \text{ km s}^{-1}]$, 50 m s^{-1} to 230 m s^{-1} above the global optimum. These solutions sample the rest of the NRHO and confirm that the global optimum at $\alpha \approx 0.44$ is a local minimum in a landscape that admits feasible transfers across most of the orbit, but with substantial cost penalty away from the pre-apolune region.
- The split between NRI and the compound mid-course manoeuvre is degenerate in the top cluster. Seeds 18 and 19/17, despite differing by only 1.7 m s^{-1} in total cost, partition the on-board budget very differently: Seed 18 has $\|\Delta v_{\text{NRI}}\| = 343 \text{ m s}^{-1}$ and a compound mid of 438 m s^{-1} , while Seeds 19/17 have $\|\Delta v_{\text{NRI}}\| \approx 470 \text{ m s}^{-1}$ and a compound mid of $\approx 314 \text{ m s}^{-1}$ with the entirety of the mid-course work concentrated on mid1 and mid2 collapsing essentially to zero (0 m s^{-1} to 2 m s^{-1}). The compound on-board ΔV is 781 m s^{-1} to 783 m s^{-1} in all three seeds.

Active constraints at convergence. At the global optimum (Seed 18), the four operational constraints behave as follows:

- Perigee altitude: $h_p = 578 \text{ km}$, within the admissible range $[150 \text{ km}, 600 \text{ km}]$ but -22 km below the upper bound. The saturation on the upper bound, universal in Cases 1–4a, is not maintained in Case 5a: the additional degree of freedom on α allows the optimiser to settle at a slightly lower perigee altitude with a marginally different LEO geometry.

- Total time of flight: $\tau_{\text{tot}} \approx 4.99$ d, marginally active at the Case 2 upper bound of 5.00 d. The TOF constraint remains marginally active, consistent with the role observed in Cases 3a–4a.
- LEO inclination: $i = 31.1^\circ$, at the upper bound of the Cape Canaveral admissible range. The constraint saturates from above, as in Cases 3a and 4a.
- NRHO insertion equality: the position and velocity at the end of segment 1 match the reference NRHO state at $\alpha = 0.4376$ to within $\|\mathbf{c}\|_\infty = 3.9 \times 10^{-5}$, well below both the operational and the strict tolerances.

Operational implication. The Case 5a result provides the first direct quantification, within the cumulative-constraint sequence of this chapter, of the cost of the apolune-insertion design choice that has been standard in the LEO $\rightarrow$ NRHO transfer literature. The marginal saving of -15 m s^{-1} relative to Case 4a is operationally insignificant for a typical crewed mission (less than 0.5% of the propellant budget), and it is recovered at the price of (i) a more complex mid-course manoeuvre profile with a compound magnitude 438 m s^{-1} rather than the 256 m s^{-1} of Case 4a, and (ii) the loss of the navigational advantage of inserting precisely at the apolune passage, where the spacecraft–Moon geometry is least sensitive to insertion errors. For a real crewed mission, the trade-off favours retention of the apolune insertion: the 15 m s^{-1} saving does not justify the operational complications. The Case 5a analysis is therefore best read as a diagnostic: it demonstrates that the apolune insertion adopted throughout Cases 1–4 is consistent with the global optimum to within 0.4% of the total mission cost, and that the cumulative-constraint sequence of this chapter has not been distorted in any operationally meaningful way by the assumption of the apolune-locking. The corresponding analysis for the Kourou launch site is reported in the following section, where the much larger plane-change penalty already documented in Case 3b is expected to interact non-trivially with the additional degree of freedom on α .

5.7 Effect of the Free NRHO Arrival Point: Kourou (Case 5b)

5.7.1 Motivation and parametrisation

Case 5b applies the same decision-space extension introduced in Section 5.5.1 - the NRHO arrival phase $\alpha \in [0, 1)$ added as a 13-th decision variable - to the Kourou launch site ($4^\circ \leq i \leq 8^\circ$). The cumulative-constraint set comprises the four operational constraints of Cases 2–4b ($\tau_{\text{tot}} \leq 5$ d, $h_p \in [150 \text{ km}, 600 \text{ km}]$, the Kourou inclination window, and the Van Allen thresholds $t_{\text{inner}} \leq 0.50$ h and $t_{\text{outer}} \leq 2.00$ h); the only change relative to Case 4b is the release of the apolune lock on the arrival point, as for Case 4a.

The qualitative expectation differs from the Cape branch. The Kourou inclination window is geometrically tight (4° wide), and the required plane change from a near-equatorial LEO to the lunar orbital plane is substantially larger. As documented in Section 5.3, this plane-change requirement adds approximately 506 m s^{-1} to the on-board

budget relative to the Cape branch, forcing heavy plane-corrections at the mid-course burns that absorb the mismatch between the LEO orbit and the target NRHO state. Releasing the arrival phase introduces a degree of freedom that is geometrically more valuable for the tightly-constrained Kourou case than for Cape, since it allows the optimiser to align the trans-lunar arc with the NRHO state at the insertion point and relieve the cislunar leg of part (or all) of the plane-correction effort. The expectation is therefore that the free arrival point will yield a larger ΔV saving in Case 5b than the -15 m s^{-1} recovered in the Cape branch (Section 5.5).

The numerical tolerances and the multi-start protocol are identical to those of Case 5a.

5.7.2 Optimal decision vector

The hybrid PSO–SQP scheme converges to the solution summarised in Table 5.21 (Seed 11), with a total mission cost $J^* = 3.842 \text{ km s}^{-1}$, a total realised flight time $\tau_{\text{tot}} = 5.000 \text{ d}$ (at the upper bound of the Case 2 constraint), and an NRHO arrival phase $\alpha^* = 0.641$ in the perilune-perilune normalisation.

Table 5.21. Optimal decision vector $\mathbf{x}^*$ for Case 5b (Kourou, $4^\circ \leq i \leq 8^\circ$, Van Allen constraint, free NRHO arrival point). Velocity components are expressed in the lunar orbital plane-aligned frame.

Variable	Description	Value	Unit
$\Delta v_{\text{NRI},x}$	NRHO insertion impulse, x	+0.0241	km/s
$\Delta v_{\text{NRI},y}$	NRHO insertion impulse, y	+0.8000	km/s
$\Delta v_{\text{NRI},z}$	NRHO insertion impulse, z	−0.0093	km/s
τ_1	Duration, segment 1	2.215	d
$\Delta v_{\text{mid1},x}$	Mid-course 1, x	+0.0012	km/s
$\Delta v_{\text{mid1},y}$	Mid-course 1, y	+0.0057	km/s
$\Delta v_{\text{mid1},z}$	Mid-course 1, z	+0.0001	km/s
τ_2	Duration, segment 2	2.065	d
$\Delta v_{\text{mid2},x}$	Mid-course 2, x	+0.0006	km/s
$\Delta v_{\text{mid2},y}$	Mid-course 2, y	+0.0003	km/s
$\Delta v_{\text{mid2},z}$	Mid-course 2, z	−0.0001	km/s
$\tau_{3,\text{max}}$	Upper bound, segment 3	8.739	d
α	NRHO arrival phase (free)	0.1461	–
<i>Derived quantities and diagnostics</i>			
$\ \Delta \mathbf{v}_{\text{NRI}}\ $	NRHO insertion magnitude	0.800	km/s
$\ \Delta \mathbf{v}_{\text{mid1}}\ $	Mid-course 1 magnitude	0.006	km/s
$\ \Delta \mathbf{v}_{\text{mid2}}\ $	Mid-course 2 magnitude	0.001	km/s
Compound mid	$\ \Delta \mathbf{v}_{\text{mid1}} + \Delta \mathbf{v}_{\text{mid2}}\ $	~ 0.007	km/s
J^*	Total mission cost	3.842	km/s
τ_{tot}	Total realised flight time	5.000	d
h_p	LEO perigee altitude	600.0	km
i	LEO inclination	8.00	deg
$d_{\text{Moon}, \text{min}}$	Minimum spacecraft–Moon distance	66 923	km
$\ \mathbf{c}\ _\infty$	Maximum constraint violation	3.9×10^{-7}	–

The capture impulse $\|\Delta \mathbf{v}_{\text{NRI}}\| = 800 \text{ m s}^{-1}$ is markedly larger than the corresponding Case 4b value of 474 m s^{-1} (Section 5.4.2), but the two mid-course magnitudes collapse to

6 m s^{-1} and 1 m s^{-1} at convergence, compared with 407 m s^{-1} and 373 m s^{-1} respectively in Case 4b. The on-board budget $\Delta V_{\text{sc}} = 807 \text{ m s}^{-1}$ is consequently 447 m s^{-1} lower than in Case 4b (1254 m s^{-1}), despite the larger NRI. The TLI cost $\Delta v_{\text{circ}} = 3035 \text{ m s}^{-1}$ is comparable to Cases 3b/4b (3046 m s^{-1} and 3012 m s^{-1} respectively).

5.7.3 Trajectory geometry

The optimised trajectory is displayed in the lunar orbital plane-aligned ECI frame in Figure 5.20. The qualitative comparison with the Case 4b apolune-locked solution, and with the Case 5a Cape branch of Section 5.5.3, reveals two features.

- The insertion point is displaced from apolune by $\Delta \text{idx}_{\text{norm}} \approx +0.14$, markedly larger than the $|\text{idx}_{\text{opt}}| = 0.0662$ found for Case 5a. Rather than settling back onto the apolune-insertion geometry of Case 4b, the Case 5b optimiser exploits the free- α freedom to displace the insertion point substantially further along the NRHO, consistent with the transition to the near-two-impulse architecture discussed above: the larger displacement is precisely what enables the mid-course corrections to collapse to near-zero, in contrast with the more modest apolune-anchored adjustment sufficient for Cape Canaveral.
- The transfer retains its direct, no-flyby character. The minimum spacecraft–Moon distance is $d_{\text{Moon, min}} = 66\,923 \text{ km}$, marginally outside the lunar sphere of influence ($\sim 66\,000 \text{ km}$) and comparable to the Case 5a value ($69\,632 \text{ km}$). No qualitatively new geometry such as a powered lunar flyby or a manifold-assisted insertion emerges from releasing the arrival phase.

The combination of these two features - the optimiser keeping the insertion near apolune yet finding a substantially lower-cost solution than Case 4b - is the central finding of Case 5b. The mechanism is not a drastic relocation of the insertion point but the architectural collapse of the on-board impulse profile, examined quantitatively in the next section.




Figure 5.20. Case 5b optimal transfer in the lunar orbital plane-aligned ECI frame, with the free NRHO arrival point active.

The Van Allen radiation transit remains operationally non-binding at the Case 5b optimum, in continuity with the behaviour already documented in all the earlier cases of the cumulative-constraint sequence. Both threshold crossings fall comfortably below the operational limits: $t_{inner} = 0.323$ h and $t_{outer} = 1.728$ h; the qualitative profile is essentially identical to that of Case 4b (Section 5.4.2), as expected from the nearly identical LEO end-state ($h_p = 600$ km, $i = 8.00^\circ$ in both cases).

5.7.4 ΔV budget and the architectural collapse of the on-board impulse profile

The total cost $J^* = 3.842 \text{ km s}^{-1}$ lies -425 m s^{-1} (-10.0%) below the Case 4b value (4.266 km s^{-1}). The saving is approximately 28 times larger than the analogous Case 5a saving relative to Case 4a (-15 m s^{-1} , -0.4%), confirming the geometric intuition of Section 5.7.1 that the Kourou launch site is substantially more sensitive to the assumption of the apolune-lock than Cape Canaveral. The detailed decomposition is reported in Table 5.22.

Two findings deserve emphasis.

The architectural collapse of the on-board impulse profile. The two mid-course corrections of the nominal three-impulse parametrisation collapse to negligible magnitudes at the converged optimum ($\|\Delta \mathbf{v}_{mid1}\| = 6 \text{ m s}^{-1}$, $\|\Delta \mathbf{v}_{mid2}\| = 1 \text{ m s}^{-1}$), and the entirety of the on-board ΔV is absorbed by a single, correspondingly larger NRHO-insertion

Table 5.22. ΔV budget for Case 5b, compared with the Kourou baseline of Case 3b (apolune insertion, no Van Allen constraint) and with Case 4b (apolune insertion, Van Allen constraint active).

Manoeuvre	Case 3b [m/s]	Case 4b [m/s]	Case 5b [m/s]	Δ (5b–4b) [m/s]
LEO insertion (TLI, Δv_{circ})	3 046	3 012	3 035	+23
Mid-course correction 1	428	407	6	−401
Mid-course correction 2	347	373	1	−372
Compound mid	775	778	~ 7	−771
NRHO insertion (Δv_{NRI})	473	474	800	+326
On-board total ΔV_{sc}	1 248	1 254	807	−447
Total J^*	4 262	4 266	3 842	−425

impulse ($\|\Delta \mathbf{v}_{\text{NRI}}\| = 800 \text{ m s}^{-1}$). The free arrival point therefore does not merely refine the geometric matching between the trans-lunar arc and the NRHO state, as it does in the Cape branch (Section 5.5.4); it enables a structural change in the impulse architecture, reducing the three-impulse design to a practically equivalent two-impulse design (TLI–NRI).

The mid-course collapse trades against a larger NRI. The net on-board saving of -447 m s^{-1} is the algebraic sum of two opposing contributions: -771 m s^{-1} from the mid-course collapse and 326 m s^{-1} from the NRI growth. The trade is energetically favourable because the cislunar mid-course corrections were operating in a region of the trans-lunar arc with a large component of out-of-plane velocity to be corrected (consistent with the heavy plane-corrective character of the Kourou apolune-locked solution documented in Section 5.3), whereas the NRI is applied at a point where the local NRHO velocity is small (the apolune region) and where the relative velocity to be matched is correspondingly small even after absorbing the plane-mismatch. The net effect is a redistribution of the on-board effort from a heavy cislunar correction to a slightly larger but more efficient terminal capture, with a substantial overall propellant saving.

The decomposition of the four impulses in the local Radial-Tangential-Normal (RTN) frame is reported in Figure 5.21, computed at the pre-impulse spacecraft state. The TLI impulse is essentially purely tangential ($T \approx 3035 \text{ m s}^{-1}$, $R, N \approx 0$), as expected for an optimal energy-injection burn at perigee. The NRI impulse is dominated by the tangential component, consistent with its primary role as a braking manoeuvre along the NRHO tangent direction at the arrival point, with smaller radial and normal contributions that account for the plane-mismatch absorbed by this single impulse. The two mid-course nodes contribute negligibly to the budget and are visible only as residual markers in the figure.




Figure 5.21. Decomposition of the four impulses of Case 5b in the local Radial-Tangential-Normal (RTN) frame, computed at the pre-impulse spacecraft state.

5.8 Sensitivity Analysis




Figure 5.22. Sensitivity analysis for Case 5b

An analogous sensitivity analysis is carried out for Case 5b (Kourou), shown in Figure 5.22. Unlike the Cape Canaveral case, the transition to the asymptotic optimum is not smooth: for bound widths up to 0.10 the PSO remains trapped in a sub-optimal family with Δv decreasing only marginally, while at $b = 0.1461$ the solution jumps discontinuously to the true global optimum ($\Delta v = 3.8413$ km/s), corresponding to a substantially larger displacement $|\text{idx}_{\text{opt}}| = 0.1461$ from the apolune than in Case 5a. This confirms that the Case 5b landscape is genuinely multimodal: a narrow admissible bound systematically excludes the global optimum, and only a sufficiently wide search window allows the optimiser to reach it.

5.8.1 Multi-start convergence and active constraints

Multi-start convergence. The twenty independent multi-start seeds yielded fourteen operationally feasible solutions ($\|\mathbf{c}\|_{\infty} \leq 10^{-3}$). The arrival-phase landscape exhibits one dominant basin and two secondary basins. Table 5.23 reports the ten lowest-cost feasible seeds. Nine of these converge to the primary basin at $\alpha \in [0.146, 0.199]$, with costs $J \in [3.842 \text{ km s}^{-1}, 3.971 \text{ km s}^{-1}]$; the tenth (Seed 1, $J = 4.048 \text{ km s}^{-1}$) already belongs to the secondary basin discussed below. Within the primary basin, the compound mid-course magnitude $\|\Delta v_{\text{mid1}} + \Delta v_{\text{mid2}}\|$ grows monotonically with rank, from a near-fully-collapsed 6 m s^{-1} (Seed 11, the case result) up to 322 m s^{-1} (Seed 14): higher-cost seeds within the same basin do not merely scale the same manoeuvre profile, but progressively reintroduce a non-trivial mid-course correction, consistent with the architectural collapse documented in Section 5.7.4 being a property of the basin optimum rather than of the basin as a whole. The two top-cost seeds (Seeds 11 and 2) converge to $J \approx 3.842 \text{ km s}^{-1}$ within 0.5 m s^{-1} of one another, both at $\alpha \approx 0.146$. Two secondary basins are visible in the multi-start (Seeds 1, 18 at $\alpha \approx -0.28$, $J \in [4.05 \text{ km s}^{-1}, 4.14 \text{ km s}^{-1}]$; Seeds 20, 8 at $\alpha \approx +0.41$, $J \in [4.13 \text{ km s}^{-1}, 4.14 \text{ km s}^{-1}]$), but with costs 200 m s^{-1} to 300 m s^{-1} above the global optimum, and they do not displace the dominance of the primary basin.

Table 5.23. Multi-start convergence of Case 5b (Kourou, Van Allen, free NRHO arrival point), reporting the ten lowest-cost seeds.

Seed	J [km/s]	$\ \mathbf{c}\ _{\infty}$	α	$\ \Delta v_{\text{NRI}}\ $ [m/s]	$\ \Delta v_{\text{mid1}}\ $ [m/s]	$\ \Delta v_{\text{mid2}}\ $ [m/s]	comp. mid [m/s]
11	3.8415	-3.9×10^{-7}	0.1461	800.4	5.8	0.7	6.3
2	3.8420	$+3.0 \times 10^{-7}$	0.1463	800.2	6.2	1.1	7.2
12	3.8554	$+3.0 \times 10^{-4}$	0.1804	635.3	184.6	1.1	185.3
6	3.8591	$+2.9 \times 10^{-5}$	0.1989	678.9	109.8	35.8	145.2
16	3.8668	-1.1×10^{-1}	0.1461	541.2	274.8	6.4	278.7
3	3.8790	-1.0×10^{-2}	0.1461	580.1	216.9	46.9	263.3
5	3.8836	$+1.2 \times 10^{-4}$	0.1749	668.9	123.8	55.9	169.8
7	3.9672	-1.5×10^{-1}	0.1461	634.1	198.5	111.7	289.0
14	3.9707	$+6.4 \times 10^{-5}$	0.1867	628.7	38.1	286.7	321.6
1	4.0483	$+8.2 \times 10^{-7}$	-0.2988	276.0	371.2	371.2	742.5

Active constraints at convergence. At the global optimum (Seed 11), the four operational constraints behave as follows: the perigee altitude is $h_p = 600.0$ km, saturating the upper bound of the admissible range as in Cases 3b/4b; the total time of flight is

$\tau_{\text{tot}} = 5.000$ d, exactly at the upper bound of the Case 2 constraint; the LEO inclination is $i = 8.00^\circ$, saturating the upper Kourou bound; the NRHO insertion equality is satisfied to within $\|\mathbf{c}\|_\infty = 3.9 \times 10^{-7}$. The Van Allen radiation transit is non-binding.

Operational implication. The Case 5b result demonstrates that releasing the apolune-insertion convention for the Kourou launch site yields a substantial ΔV saving (-425 m s^{-1} , -10.0%) by enabling an architectural transition from a heavy three-impulse cislunar architecture (Case 4b) to an essentially two-impulse TLI–NRI design with all on-board work concentrated at the terminal capture. The saving is approximately 84% of the Kourou plane-change penalty (Case 3b relative to the Case 1 baseline, $+506 \text{ m s}^{-1}$) and is large enough to be operationally relevant for a real mission design: it directly reduces the propellant mass at launch by a comparable percentage. The trade-off, however, is that the architectural collapse concentrates the on-board braking authority at a single 800 m s^{-1} impulse, which couples the entire mission feasibility to the successful execution of that manoeuvre and forecloses the navigational role of the mid-course corrections as a real-time error-recovery mechanism. Whether this architectural simplification is operationally acceptable depends on considerations external to the optimisation framework of this thesis (propulsion system sizing, navigational redundancy, contingency planning) and is not pursued further here. For the cumulative-constraint analysis of this chapter, the conclusion is that the apolune-insertion convention adopted throughout Cases 1–4 is operationally near-optimal for the Cape branch (within 15 m s^{-1}) but is suboptimal by $\sim 425 \text{ m s}^{-1}$ for the Kourou branch. The sensitivity of this conclusion to the choice of NRHO arrival epoch within the launch window is the subject of the next section.

5.9 Secular Study: Cape Canaveral versus Kourou

The Case 5a and 5b analyses of Sections 5.5 and 5.7 have so far considered a single reference NRHO arrival epoch (2028-03-15). Because of the strong precession of the NRHO across consecutive synodic periods, the geometric matching between the lunar orbital plane and the LEO launch site - and consequently the cost of the LEO $\rightarrow$ NRHO transfer — depends explicitly on the choice of arrival epoch within the launch window. This section quantifies that dependence by repeating the Case 5 multi-start optimisation across twenty consecutive NRHO periods, for both the Cape Canaveral and the Kourou launch sites.

Methodological scope. The secular study uses, for each NRHO period $k \in \{0, \dots, 19\}$, an apolune-anchored parametrisation: the optimiser’s normalisation parameter is set to $P.\text{idx_norm} = \alpha_{\text{apo}}(k)$, with $x_{13} \in [0, 1]$ in the perilune-to-perilune normalisation representing the signed displacement of the arrival phase from the local apolune in fractions of a period. However, as documented in Section 5.7.2 for the reference epoch, the Case 5b cost landscape admits structurally distinct architectural basins, only one of which is fully accessible from the apolune-anchored multi-start; the implications of this for the interpretation of the secular results are discussed at the end of the section.

The three-panel summary in Figure 5.23 reports, for each NRHO period, *(i)* the optimal ΔV , *(ii)* the optimal arrival phase α^* in the perilune-to-perilune coordinate, and *(iii)* the corresponding physical time offset of the insertion point relative to apolune. The contrast between the two launch sites is qualitative rather than merely quantitative.




Figure 5.23. Secular variation of the Case 5 free-arrival-point optima over twenty consecutive NRHO periods.

The findings are summarised in Table 5.24 and discussed in the paragraphs that follow.

The Cape Canaveral optimum is structurally stable across the launch window.

For Case 5a, the optimal arrival phase remains in a narrow band around apolune ($\sigma_{\alpha^*} = 0.07$, $\sigma_{\Delta t_{\text{apo}}} = 11$ h) across all sixteen feasible NRHO periods, and the associated ΔV dispersion is correspondingly small ($\sigma_{J^*} = 27 \text{ m s}^{-1}$, total range 85 m s^{-1}). The cost landscape of Case 5a admits a single dominant basin of attraction, centred at or near apolune; the multi-start across NRHO periods samples this basin consistently, with only second-order modulations due to the precession of the orbit within the synodic month. The single-basin character of the Cape 5a landscape implies that the apolune-anchored secular sampling adopted in this section captures the actual global optimum at each launch epoch, with the residual dispersion faithfully representing the true sensitivity of the global optimum to the launch-epoch geometry.

Table 5.24. Secular variation statistics over twenty NRHO periods of the Artemis IV launch window.

Quantity	Case 5a (Cape)	Case 5b (Kourou)
N_{feas}	16/20	9/20
Mean J^* [km/s]	3.785	3.878
Std. dev. J^* [m/s]	27	120
Range J^* [km/s]	[3.735, 3.821]	[3.751, 4.218]
Range width J^* [m/s]	85	470
Mean α^*	0.475	0.363
Std. dev. α^*	0.072	0.322
Range α^*	[0.39, 0.67]	[−0.08, 0.86]
Mean Δt_{apo} [h]	−0.9	+5.5
Std. dev. Δt_{apo} [h]	11.4	49.0
Range Δt_{apo} [h]	[−17, +28]	[−47, +69]
Reference epoch p_0 baseline J^* [km/s]	3.821	3.842

The Kourou optimum is multi-basin and substantially more dispersed. For Case 5b, the optimal arrival phase scatters over almost the full orbital period ($\sigma_{\alpha^*} = 0.32$, $\sigma_{\Delta t_{\text{apo}}} = 49$ h, range $\alpha^* \in [-0.08, 0.86]$), and the ΔV dispersion is one order of magnitude larger than for Case 5a ($\sigma_{J^*} = 120$ m s^{−1}, total range 470 m s^{−1}). The landscape admits multiple feasible basins (near-apolune $\alpha \approx 0.5$, pre-apolune $\alpha \approx 0.2$, and post-apolune $\alpha \approx 0.85$, corresponding respectively to apolune-adjacent insertion, insertion near the inbound transition, and insertion near the outbound transition to the next perilune); which basin contains the optimum depends on the specific geometry of the departure epoch within the NRHO precession cycle.

Physical interpretation. The qualitative difference between the two launch sites is best read as a consequence of the geometric coupling between the launch-site inclination window and the required plane change. For Cape Canaveral, the relatively wide window ($26^\circ \leq i \leq 31^\circ$, nominal NRHO plane match achievable with moderate effort) allows the optimiser to consistently access the apolune-adjacent basin, which is energetically preferred because the local NRHO velocity is minimal and the mismatch with the trans-lunar arc is smallest at the orbital extremum. For Kourou, the much narrower window ($4^\circ \leq i \leq 8^\circ$, requiring a much larger plane change to reach the NRHO plane) opens additional architectural basins (with redistribution of the plane-correction effort between cislunar mid-course corrections and the terminal NRI capture), and the optimal architecture at any given launch epoch depends on the specific orientation of the LEO launch plane and the NRHO plane at the candidate arrival epoch. The multimodality of Case 5b is therefore a direct manifestation of the tighter feasible domain of the Kourou problem, both in α^* location and in ΔV architecture; it is not a sampling artefact of the multi-start.

Operational implication. The contrast between the two launch sites has a direct consequence for the mission design phase.

For a Cape Canaveral launch, the apolune-insertion design choice is robust across the launch window: the cost of any specific NRHO arrival epoch in the operational window lies within a 85 m s^{−1} envelope of the global optimum, and the propellant budget can be sized on a single nominal value with a correspondingly small contingency.

For a Kourou launch, the propellant budget cannot be determined directly from the secular dispersion of Table 5.24, because two operational scenarios bracket the actual budget. In the optimistic scenario, the apolune-adjacent architecture is accessible across the entire launch window with performance similar to the reference epoch ($\sim 3.84 \text{ km s}^{-1}$); in this case the Kourou budget is competitive with Cape (within $\sim 60 \text{ m s}^{-1}$). In the pessimistic scenario, the apolune-adjacent architecture is not consistently accessible, and at unfavourable epochs the mission would have to fall back to the three-impulse-with-SOI-passage architecture sampled by the secular study, with ΔV up to 4.22 km s^{-1} at the worst-case epoch in the launch window. The propellant contingency to cover the pessimistic scenario is 300 m s^{-1} to 500 m s^{-1} above the optimistic-scenario value, depending on the worst-case epoch retained.

5.10 Cumulative Summary

The cumulative-constraint sequence analysed throughout this chapter is consolidated in Table 5.25, which collects the converged cost of the best operationally feasible seed of each case together with its marginal cost relative to the preceding case of the same launch-site branch. Two features of the parametrisation make the table readable as a single cost narrative. First, the trans-lunar injection Δv_{circ} remains within a narrow band of approximately $3.01\text{--}3.08 \text{ km/s}$ across every case, set almost entirely by the energy required to leave the LEO parking orbit; the operational constraints redistribute the on-board budget ΔV_{sc} (the NRHO insertion and the two mid-course corrections) rather than the injection itself, and it is therefore ΔV_{sc} that carries the signature of each restriction. Second, the two launch-site branches diverge from Case 3 onward and are best read in parallel.

Table 5.25. Cumulative summary of the LEO→NRHO cases.

Case	New element	J^* [km/s]	$\Delta(\Delta V)$ [m/s]	ΔV_{sc} [m/s]
1	Perigee altitude only	3.783	—	742
2	+ TOF $\leq 5 \text{ d}$	3.782	−1	741
<i>Cape Canaveral branch</i>				
3a	+ inclination $26\text{--}31^\circ$	3.829	+47	789
4a	+ Van Allen (operational)	3.836	+7	796
4c	Van Allen (tightened) [†]	3.866	+30	810
5a	+ free arrival point α	3.821	−15	781
<i>Kourou branch</i>				
3b	+ inclination $4\text{--}8^\circ$	4.261	+479	1248
4b	+ Van Allen (operational)	4.266	+5	1254
4d	Van Allen (tightened) [†]	4.394	+128	1313
5b	+ free arrival point α	3.842	−425	807

[†] Marginal cost of Cases 4c/4d is referred to the operational Cases 4a/4b; that of Cases 5a/5b to Cases 4a/4b; that of Cases 3a/3b to Case 2.

Time of flight is non-binding. The transition from Case 1 to Case 2 leaves the optimum essentially unchanged (-1 m/s in cost, -0.02 d in flight time): the five-day time-of-flight limit excludes the long-flight-time basin from the feasible domain but does not raise the cost of the short-flight-time optimum, which already satisfies the bound with margin. The time-of-flight constraint is therefore inactive at every converged optimum of the chapter.

Launch-site inclination is the dominant cost driver. The inclination constraint is the single most expensive restriction of the cumulative chain, and it exposes a pronounced asymmetry between the two launch sites. The Cape Canaveral window ($26\text{--}31^\circ$) costs only $+47$ m/s of on-board budget, whereas the Kourou window ($4\text{--}8^\circ$) costs $+506$ m/s relative to the unconstrained baseline—approximately eleven times as much—despite the plane misalignment with the NRHO being only about 3.3 times larger. This super-linear scaling is the quantitative signature of the geometric coupling between the launch site and the NRHO: the Cape Canaveral inclination band aligns naturally with the lunar orbital plane, while the near-equatorial Kourou band forces a costly plane change that the optimiser absorbs into a compound mid-course manoeuvre.

Van Allen transit is non-binding at the operational thresholds. The operational belt limits derived from Artemis I mission (Cases 4a/4b) are satisfied with margin at the inclination-constrained optima, and their activation changes the cost by only $+7$ m/s (Cape) and $+5$ m/s (Kourou), differences attributable to an alternative organisation of the mid-course nodes rather than to a structurally different solution. A measurable cost appears only under the tightened-belt sensitivity variants (Cases 4c/4d, $+30$ and $+128$ m/s), and even there the reported figures are lower bounds: the converged optima retain a small residual outer-belt excess (6.8 and 7.5 min against the 1.50 h tightened threshold), so the true cost of fully meeting the tightened limit would be higher.

The free arrival point. Releasing the apolune-insertion convention has a qualitatively different effect on the two branches. For Cape Canaveral (Case 5a) the saving is a negligible -15 m/s: the apolune is already close to the cost-optimal arrival point, and the additional degree of freedom only refines it. For Kourou (Case 5b) the same relaxation yields -425 m/s. The mechanism is an architectural transition rather than a local refinement: the compound mid-course manoeuvre collapses from 778 m/s to a few m/s, and the entire on-board budget concentrates in a single, correspondingly larger NRHO-insertion impulse of ~ 800 m/s, producing an essentially two-impulse TLI–NRI transfer. This recovers approximately 84% of the Kourou plane-change penalty by allowing the arrival phase to absorb the plane mismatch into the braking impulse rather than paying for it with a dedicated mid-course.

Net cumulative envelope. The combined effect of the full operational constraint set and the free arrival point is that both branches return close to the unconstrained baseline of 3.783 km/s. The Cape Canaveral chain terminates at 3.821 km/s (Case 5a), only $+38$ m/s above the baseline, and the Kourou chain at 3.842 km/s (Case 5b), only $+59$ m/s above it. The free-arrival relaxation thus very nearly cancels the cumulative penalty of the operational constraints: the two launch sites, which differ by some 430 m/s once the inclination constraint and the convention of the apolune lock are both enforced (Cases 4a vs. 4b), converge to within 21 m/s of one another once the arrival point is freed. The

apolune-insertion convention adopted throughout Cases 1–4 is consequently near-optimal for Cape Canaveral but materially suboptimal for Kourou.

Parametric degeneracy and epoch sensitivity. A methodological observation recurs across the constrained cases: where the plane change is large, the true decision variable at convergence is the compound mid-course vector, not its split between the two individual nodes. Finally, the cost figures of Table 5.25 are reported at a single reference epoch; their sensitivity to the NRHO arrival epoch within the launch window is the subject of the secular-variation study of Section 5.9, which shows the Cape Canaveral optimum to be structurally stable across consecutive synodic periods and the Kourou optimum to be multi-basin and substantially more dispersed.

6 | Scenario 2: NRHO to Powered Descent Initiation

This chapter presents the optimisation results for Scenario 2, the descent transfer from the southern L_2 NRHO to the Powered Descent Initiation (PDI) gate at an altitude of $h_{\text{PDI}} = 15$ km above the mean lunar surface. The methodology, dynamical model and constraint structure were established in Section 4.4; the present chapter applies them to the five cumulative cases enumerated in Section 4.4.3 and to the corresponding secular-variation study.

The exposition mirrors the structure adopted for Scenario 1: each case is presented through its optimal decision vector, trajectory geometry in the selenocentric inertial frame, ΔV budget, and multi-start convergence behaviour. The cumulative logic is preserved: each case retains the active constraint set of the preceding one and introduces a single new requirement, isolating the marginal ΔV impact of each operational restriction. Case 1 (Section 6.2) adopts a reduced single-impulse parametrisation as a diagnostic baseline; Cases 2–6 share the three-impulse parametrisation of Eq. (4.14), with Case 6 extending the decision vector to admit a free departure point on the NRHO.

A preliminary remark concerns the PDI velocity constraint. The methodology introduces a two-sided inequality on the Moon-relative speed at the PDI gate, $|v_{\text{PDI}}| \in [1.660, 1.710]$ km/s about the nominal value of 1.685 km/s, derived from the entry conditions of the powered-descent guidance loop. A first analysis run revealed that the unconstrained ΔV minimum of the three-impulse parametrisation falls outside this band: the natural arrival speed at $h = 15$ km, with no upper-velocity penalty, is approximately 1.77 km/s. The strict constraint can be satisfied at the expense of additional propellant expenditure, but the resulting trade-off between $|v_{\text{PDI}}|$ and ΔV is not unique to a single operating point: along the descent geometry, configurations that substantially reduce ΔV admit arrival speeds up to approximately 2.1 km/s. In the present analysis the constraint is therefore interpreted as a soft operational specification: convergence is declared when the maximum residual of the active constraint set satisfies the operational threshold $\max_k c_k(\mathbf{x}^*) \leq 0.5$, with ΔV minimisation as the primary criterion. Each reported solution is verified a posteriori for compliance with the altitude, latitude, surface-impact and (where active) TOF constraints; arrival velocities exceeding the nominal band are documented case by case in the caption tables and discussed in a final cumulative summary.

6.1 Reference Configuration

All cases of this chapter share a common reference epoch, departure geometry and dynamical setup, summarised here once to avoid repetition in the case-specific sections that follow.

Epoch and ephemeris. The reference epoch for all simulations is $t_{\text{ref}} = 10$ March 2028, 00:00:00 UTC. The N -body dynamics is evaluated against the JPL DE440 ephemeris loaded through the SPICE/MICE toolkit; the spacecraft’s reference NRHO state is provided by the kernel `receding_horiz_3189_1burnApo_DiffCorr_15yr.bsp` (SPK identifier `-60000`), which spans 15 years of the NASA nine-by-two synodic-resonant southern L_2 orbit. The body-fixed lunar frame `IAU_MOON`, required for the evaluation of the selenographic latitude at the PDI gate, is loaded through the planetary constants kernel `pck00011.tpc`.

Departure state. The fully-specified anchor of each propagation is the NRHO perilune state nearest to the reference epoch. Numerically, the perilune is located by scanning the spacecraft’s selenocentric distance over one full NRHO period and selecting the minimum-distance epoch; for the reference epoch adopted in this chapter the perilune phase relative to t_{ref} is $\text{idx}_{\text{norm}} = 0.3596$ (corresponding to $\Delta t \approx 2.36$ d ahead of t_{ref}), with a perilune altitude of $h_{\text{peri}} = 1638$ km above the mean lunar surface. The state vector at perilune, $\mathbf{x}_{\text{NRHO}}(t_{\text{peri}}) = [\mathbf{r}_{\text{NRHO}}, \mathbf{v}_{\text{NRHO}}]^\top$, is extracted directly from the SPICE ephemeris and is identical across all cases except Case 6, in which a free phase offset α along the NRHO is admitted as a thirteenth optimisation variable.

Selenocentric reference frame. All trajectory plots in this chapter are referred to a right-handed selenocentric inertial frame $\{\hat{X}_{\text{sel}}, \hat{Y}_{\text{sel}}, \hat{Z}_{\text{sel}}\}$ defined at the reference perilune epoch $t_{\text{peri,ref}}$, with $\hat{X}_{\text{sel}}$ along the Earth–Moon line and positive in the Earth-to-Moon outward direction (i.e. pointing away from the Earth as seen from the Moon), $\hat{Z}_{\text{sel}}$ aligned with the Moon’s instantaneous orbital angular-momentum vector $\mathbf{r}_{\mathcal{M}} \times \mathbf{v}_{\mathcal{M}}$ (positive in the sense of the lunar orbital revolution), and $\hat{Y}_{\text{sel}} = \hat{Z}_{\text{sel}} \times \hat{X}_{\text{sel}}$ completing the right-handed triad along the prograde direction of the lunar orbital motion. The rotation matrix $\mathbf{R}_{\text{sel}} = [\hat{X}_{\text{sel}}, \hat{Y}_{\text{sel}}, \hat{Z}_{\text{sel}}]^\top$ is computed once at initialisation and held fixed throughout the optimisation run, including the Case 5 free- α configuration in which the departure epoch t_{dep} varies with the decision variable α . The optimisation variables $\Delta \mathbf{v}_k$ are the impulse components expressed in this frame; the inverse rotation $\mathbf{R}_{\text{sel}}^\top \Delta \mathbf{v}_k$ is applied before each impulse is inserted into the J2000-frame propagator. The selenographic latitude at the PDI gate, φ_{PDI} , is evaluated instead in the body-fixed lunar frame `IAU_MOON` (Section 4.2.4) via the transformation $\mathbf{T}_{\text{J2000} \rightarrow \text{IAU_MOON}}(\mathbf{e}_{\text{tPDI}})$ provided by SPICE.

Numerical setup. The four-body equations of motion (Earth, Moon, Sun, with the spacecraft as a massless point particle) are integrated using MATLAB’s `ode113` adaptive Adams–Bashforth–Moulton solver, with absolute and relative tolerances of 10^{-9} and additional options `InitialStep` = 60 s and `MaxStep` = 600 s to ensure accurate triggering of the PDI altitude-gate event. The PDI gate is enforced as a terminal event on the descending crossing of the $h = 15$ km altitude surface; a complementary terminal event at $h = 0$ km guards against lunar-surface impact during the descent segment. The optimisation is performed in two stages: a particle-swarm global search with swarm size

80 and maximum 500 iterations (no early-stopping), followed by a sequential-quadratic-programming local refinement (`fmincon`, algorithm `sqp`) with default tolerances. The multi-start procedure repeats this two-stage optimisation with $N_{\text{seeds}} = 10$ independent random initialisations for Cases 1 to 4, and $N_{\text{seeds}} = 50$ for the free-departure Case 5, where the 13-dimensional decision space and the observed multimodality of the solution landscape justify the additional sampling.

6.2 Case 1: Single-Impulse Baseline

The first case adopts a reduced parametrisation in which a single impulse $\Delta \mathbf{v}_{\text{dep}}$ is applied at the NRHO perilune and the spacecraft coasts ballistically to the PDI gate. The decision vector is therefore four-dimensional,

$$\mathbf{x}^{(1)} = [\Delta v_{\text{dep},x}, \Delta v_{\text{dep},y}, \Delta v_{\text{dep},z}, \tau_1]^\top \in \mathbb{R}^4, \quad (6.1)$$

with the three impulse components expressed in the selenocentric frame and the time of flight τ_1 bounded in $[0.05, 0.60]$ d. The only active constraints are the PDI altitude window, the PDI velocity window, and the lunar-surface impact safeguard; no geometric or temporal restrictions on the arrival point are imposed.

The role of this case is diagnostic: by removing the mid-course and terminal-braking impulses from the parametrisation, the resulting ΔV^* establishes a quantitative upper bound on the cost of the descent transfer. The subsequent introduction of the three-impulse parametrisation in Case 2 (Section 6.3) can then be assessed against this baseline through a direct ΔV comparison.

6.2.1 Optimal Decision Vector

The lowest- ΔV seed satisfying the operational tolerance $\max_k c_k(\mathbf{x}^*) \leq 0.5$ is reported in Table 6.1. The converged impulse magnitude is $|\Delta \mathbf{v}_{\text{dep}}^*| = 1.0021$ km/s, with the spacecraft reaching the PDI gate after a coast of $\tau_1 = 0.034$ d.

Table 6.1. Converged decision vector for Case 1 (single-impulse baseline), seed 1 of 10. Components expressed in the selenocentric inertial frame $\{X_{\text{sel}}, Y_{\text{sel}}, Z_{\text{sel}}\}$ of Section 6.1.

Component	Value	Unit
$\Delta v_{\text{dep},x}$	+0.0088	km/s
$\Delta v_{\text{dep},y}$	-1.0000	km/s
$\Delta v_{\text{dep},z}$	-0.0644	km/s
$ \Delta \mathbf{v}_{\text{dep}} $	1.0021	km/s
τ_1	0.034	d

The short coast time and the magnitude of the departure impulse indicate that the trajectory is essentially a forced descent: the perilune state of the NRHO carries the spacecraft to a minimum altitude of 1638 km, significantly higher than the PDI gate; the single impulse must therefore both reduce the selenocentric energy enough to lower the

periselene to the PDI altitude band and shape the trajectory's energy distribution so that the arrival speed remains within the operational window.

6.2.2 Trajectory Geometry

The converged trajectory in the selenocentric inertial frame is shown in Figure 6.1. The departure impulse at the NRHO perilune bends the NRHO arc into a direct descent to the PDI gate. The resulting descent is essentially meridional, the departure impulse driving the spacecraft directly towards the PDI gate.




Figure 6.1. Case 1 (single-impulse) optimised trajectory in the selenocentric inertial frame.

The arrival point lies in the northern selenographic hemisphere at $\phi_{\text{PDI}} = +38.59^\circ$; no latitude constraint is active in this case, so the converged geometry reflects the unconstrained natural descent corridor accessible from the reference perilune. The selenographic latitude evolution along the descent and the corresponding altitude profile are shown in Figure 6.2. The altitude decreases monotonically from 1638 km at perilune to the 15 km PDI gate event over 0.034 d, while the latitude transitions from $\phi_0 \approx +90^\circ$ (perilune position) to $+38.6^\circ$ at PDI.




Figure 6.2. Selenographic latitude and altitude profiles for the Case 1 converged trajectory.

6.2.3 ΔV Budget

The case-1 ΔV budget is informative as a baseline, as the entire mission cost is concentrated in the single departure impulse,

$$\Delta V_{\text{tot}}^{(1)} = |\Delta \mathbf{v}_{\text{dep}}^*| = 1.0021 \text{ km/s.} \quad (6.2)$$

As presented in the following section, this value is to be compared with the three-impulse minimum of Case 2 (Section 6.3), which converges to $\Delta V^{(2)} \approx 0.742 \text{ km/s}$ under the same active-constraint set: the additional two impulses introduced by the three-impulse parametrisation reduce the mission cost by approximately 260 m/s, quantifying the saving achievable through mid-course correction and terminal braking before the PDI gate.

6.2.4 Multi-start Convergence

The ten independent PSO+SQP runs converge to a remarkably consistent optimum, as summarised in Table 6.2: all seeds fall within the band $\Delta V \in [1.0021, 1.0031] \text{ km/s}$ with identical operational violation $\max_k c_k = 6.35 \times 10^{-2}$. The constraint residual is dominated by the velocity term, $|v_{\text{PDI}}| = 1.77 \text{ km/s}$ exceeding the nominal upper bound of 1.710 km/s by $\approx 60 \text{ m/s}$; this is characteristic of the single-impulse parametrisation, which lacks the degree of freedom required to simultaneously satisfy the altitude and velocity windows.

The tightness of the multi-start solutions confirms that the four-dimensional problem of Case 1 admits a single basin of attraction in the feasible region; the global-minimum identification is robust against the stochastic initialisation.

Table 6.2. Multi-start convergence for Case 1: ΔV^* and PSO runtime per seed.

Seed	ΔV^* [km/s]	$\max_k c_k$	t_{PSO} [min]
1	1.0021	6.35×10^{-2}	1.4
2	1.0025	6.35×10^{-2}	0.5
3	1.0026	6.35×10^{-2}	0.3
4	1.0026	6.35×10^{-2}	0.3
5	1.0021	6.35×10^{-2}	0.3
6	1.0027	6.35×10^{-2}	0.3
7	1.0027	6.35×10^{-2}	0.3
8	1.0031	6.35×10^{-2}	0.3
9	1.0029	6.35×10^{-2}	0.3
10	1.0028	6.35×10^{-2}	0.3
<i>Best:</i> seed 1, $\Delta V^* = 1.0021$ km/s, $\tau^* = 0.0345$ d			

6.3 Case 2: Three-Impulse Parametrisation, Natural TOF

Case 2 introduces the three-impulse parametrisation that will be retained throughout the remaining cases. The decision vector has now twelve dimensions,

$$\mathbf{x}^{(2)} = [\Delta \mathbf{v}_{\text{dep}}^\top, \tau_1, \Delta \mathbf{v}_{\text{mcc}}^\top, \tau_2, \Delta \mathbf{v}_{\text{PDI}}^\top, \tau_3]^\top \in \mathbb{R}^{12}, \quad (6.3)$$

with three impulses ($\Delta \mathbf{v}_{\text{dep}}$ at the NRHO perilune, a mid-course correction $\Delta \mathbf{v}_{\text{mcc}}$, and a terminal braking $\Delta \mathbf{v}_{\text{PDI}}$) and three coast segments of duration τ_1, τ_2, τ_3 . Both impulse components and segment durations are constrained by the box bounds of Table 4.3; no geometric or temporal restriction on the arrival point is imposed.

The role of this case is twofold. First, it quantifies the ΔV saving achievable by splitting the descent manoeuvre into three impulses, to be compared against the single-impulse baseline of Section 6.2. Second, by leaving the total time of flight unconstrained, it exposes the “natural” descent time and arrival geometry of the three-impulse parametrisation, against which the operational constraints introduced from Case 3 onwards can be benchmarked.

6.3.1 Optimal Decision Vector

The lowest- ΔV seed satisfying the operational tolerance $\max_k c_k \leq 0.5$ is reported in Table 6.3. The converged total cost is $\Delta V^* = 0.7419$ km/s, distributed across three impulses applied over an optimised time budget of $\tau_{\text{tot}}^* = 0.754$ d. The actual flight time, set by the PDI gate event during the third segment, is $\tau_{\text{actual}} = 0.686$ d (≈ 16.5 h); the difference of 1.6 h corresponds to the unused fraction of the τ_3 budget that the truncating altitude event renders ineffective.

The departure impulse retains the predominantly retrograde-in-plane orientation observed in Case 1: $-Y_{\text{sel}}$ is again the dominant component (-0.5595 km/s, 99% of the impulse

Table 6.3. Converged decision vector for Case 2 (three-impulse, natural TOF), seed 1 of 10. Components expressed in the selenocentric inertial frame of Section 6.1.

	Departure ($k = 1$)	Mid-course ($k = 2$)	PDI braking ($k = 3$)
$\Delta v_{k,x}$ [km/s]	+0.0609	+0.0313	+0.0070
$\Delta v_{k,y}$ [km/s]	−0.5595	−0.1079	+0.0251
$\Delta v_{k,z}$ [km/s]	+0.0699	−0.0967	+0.0055
$ \Delta \mathbf{v}_k $ [km/s]	0.5671	0.1482	0.0267
τ_k [d]	0.4864	0.0827	0.1844

magnitude). The magnitude itself is roughly halved compared to Case 1 (0.567 km/s versus 1.002 km/s), confirming that the introduction of two additional correction impulses allows the departure burn to operate without saturated bounds: no component of any of the three impulses sits on its ± 1.0 km/s limit in this case. The mid-course correction follows a similar retrograde orientation in Y_{sel} with an out-of-plane Z component of comparable magnitude, while the PDI braking impulse, with $|\Delta \mathbf{v}_{\text{PDI}}^*| = 27$ m/s, is reduced to a small mid-air targeting correction.

6.3.2 Trajectory Geometry

The converged trajectory in the selenocentric inertial frame is shown in Figure 6.3, with the three impulse vectors made visible at their respective application points along the descent. The selenographic latitude and altitude profiles are shown in Figure 6.4.

The descent extends over approximately 16.5 h, more than 20 times longer than the Case 1 forced descent of Section 6.2. The increased coast time admits a more energetically efficient transfer geometry: rather than driving the spacecraft directly towards the lunar surface, the trajectory follows an extended arc on which the energy is shed progressively through three discrete impulses. As is apparent in Figure 6.3, the spacecraft completes approximately two and a half revolutions about the Moon before reaching the PDI gate, with the trajectory’s farthest excursion from the X_{sel} axis reaching ~ 2700 km in the $-Y_{\text{sel}}$ direction.

A notable feature of the converged geometry concerns the arrival latitude. With no geographic constraint imposed, Case 2 nevertheless converges to an arrival in the southern selenographic hemisphere at $\phi_{\text{PDI}} = -66.4^\circ$. This natural southern-hemisphere convergence is consequential for the cumulative-constraint sequence: the hemispheric constraint introduced in Case 3 ($\phi_{\text{PDI}} \leq 0$) is automatically satisfied by the geometry of Case 2, suggesting that the marginal ΔV cost of that constraint should be small. The -85° constraint of Case 4, in contrast, requires the latitude to extend approximately 19° further south than the natural convergence point of Case 2, and a quantifiable cost is expected.

6.3.3 ΔV Budget

The total mission ΔV for the optimal seed is distributed across the three impulses as

$$\Delta V_{\text{tot}}^{(2)} = |\Delta \mathbf{v}_{\text{dep}}^*| + |\Delta \mathbf{v}_{\text{mcc}}^*| + |\Delta \mathbf{v}_{\text{PDI}}^*| = 0.5671 + 0.1482 + 0.0267 = 0.7420 \text{ km/s}, \quad (6.4)$$

NRHO $\rightarrow$ PDI - case 2, seed 1




Figure 6.3. Case 2 (three-impulse, natural TOF) optimised trajectory in the selenocentric inertial frame.

Comparing with the single-impulse baseline of Case 1 ($\Delta V^{(1)} = 1.0021$ km/s, Section 6.2), the admission of two additional impulses reduces the mission cost by 260 m/s, a saving of 26% relative to the baseline.

The distribution of the budget across the three impulses is highly uneven: the departure impulse dominates (76.4% of ΔV_{tot}), the mid-course refinement contributes 20.0%, and the terminal braking $\Delta \mathbf{v}_{\text{PDI}}^*$ is reduced to a near-zero correction (3.6%, 27 m/s). This degenerate use of the third impulse is a direct consequence of the operational velocity tolerance accepted in this analysis: the converged seed admits an arrival speed at the PDI gate of $|v_{\text{PDI}}| = 1.8805$ km/s, exceeding the nominal upper bound of 1.710 km/s by 170 m/s. This violation accounts entirely for the residual $\max_k c_k = 1.70 \times 10^{-1}$ reported in Table 6.4 (the altitude, latitude, surface and TOF residuals are all strictly negative). The braking impulse is consequently not required to enforce the velocity window, and reduces to a small mid-air targeting correction. This trade-off, intrinsic to the natural minimum- ΔV corridor of the descent geometry, will be progressively suppressed by the




Figure 6.4. Selenographic latitude and altitude profiles. The vertical dashed lines indicate the application times of the mid-course (τ_1) and braking ($\tau_1 + \tau_2$) impulses.

cumulative constraints introduced in the subsequent cases: as the velocity-relaxed low- ΔV cluster becomes infeasible under tighter geographic restrictions, the braking impulse will recover its nominal role of bringing the spacecraft within the prescribed PDI velocity window.

6.3.4 Multi-start Convergence

In contrast to the unimodal landscape of Case 1, the three-impulse parametrisation exposes a more structured solution space: the ten multi-start seeds converge into two distinct clusters, summarised in Table 6.4. A first cluster (seeds 1 and 6) collects two low- ΔV solutions at $\Delta V \approx 0.74$ km/s with operational violation $\text{viol} \approx 0.1\text{--}0.2$ (the PDI velocity exceeds the nominal upper bound by $\sim 100\text{--}170$ m/s). A second cluster (seeds 2–5, 7–10) collects solutions at $\Delta V \approx 0.83\text{--}0.89$ km/s that comply with the velocity window ($\text{viol} \leq 10^{-1}$). The two clusters quantify the ΔV cost of enforcing the nominal velocity band rigidly: relaxing the velocity tolerance by ~ 100 m/s reduces the mission cost by ~ 100 m/s, evidence that in the three-impulse space the $\Delta V\text{--}|v_{\text{PDI}}|$ trade-off is approximately linear over the relaxation range considered.

The two-cluster convergence structure is representative of the descent transfer problem and will be observed consistently throughout the cumulative-constraint sequence: as additional constraints are activated in Cases 3–5, the velocity-relaxing low- ΔV cluster progressively disappears as the nominal-velocity cluster becomes the only feasible region.

Table 6.4. Multi-start convergence for Case 2: ΔV^* and operational violation per seed.

Seed	ΔV^* [km/s]	$\max_k c_k$	t_{PSO} [min]
1	0.7419	1.70×10^{-1}	2.2
2	0.8754	9.4×10^{-3}	1.1
3	0.8365	5.3×10^{-2}	1.0
4	0.8821	3.5×10^{-3}	0.7
5	0.8267	7.5×10^{-2}	1.2
6	0.7447	1.1×10^{-1}	0.8
7	0.8861	2.5×10^{-3}	0.9
8	0.8893	8.9×10^{-3}	1.1
9	0.8943	7.6×10^{-5}	0.7
10	0.8934	$\leq 10^{-6}$ (nominal feasible)	1.1
<i>Best</i> ($OP_TOL = 0.5$): seed 1, $\Delta V^* = 0.7419$ km/s, $\tau_{\text{actual}} = 0.686$ d			

6.4 Case 3: Southern Hemisphere Arrival, Natural TOF

Case 3 introduces the first geographic constraint of the cumulative sequence: the arrival point is restricted to the southern selenographic hemisphere through the inequality $\phi_{\text{PDI}} \leq 0^\circ$. The decision vector retains the twelve-dimensional three-impulse parametrisation of Equation (6.3); the constraint set of Case 2 is augmented by the latitude residual $c_\phi(\mathbf{x}) = \phi_{\text{PDI}}(\mathbf{x})$ alongside the active altitude, velocity and surface-impact terms.

A noteworthy feature of this case is that the new constraint is structurally inactive at the optimum of the preceding case: as stated before the lowest- ΔV seed of Case 2 (Section 6.3) arrives naturally at $\phi_{\text{PDI}} = -66.4^\circ$, well inside the admissible region of Case 3. The hemispheric constraint therefore does not in principle exclude the Case 2 optimum from the feasible set. Any ΔV increase observed between the two cases must accordingly be attributed not to constraint activation at the optimum, but to the restriction of the PSO search space that the constraint imposes during the global exploration phase: portions of the decision domain that the unrestricted Case 2 search could exploit are now masked off by the latitude penalty term, with measurable consequences for the multi-start distribution.

The lowest- ΔV seed satisfying $\max_k c_k \leq 0.5$ converges to $\Delta V^* = 0.7685$ km/s (seed 2), representing a marginal increase of ~ 27 m/s (+3.6%) over the natural-TOF optimum of Case 2. The full multi-start convergence is reported in Table 6.5; the two-cluster structure observed in Case 2 persists, with seeds 2 and 5 forming the velocity-relaxing low- ΔV cluster and the remaining seeds collecting nominal solutions with $\Delta V \approx 0.84$ – 0.91 km/s.

Two qualitative observations frame the result. First, the converged arrival latitude of the optimal seed, $\varphi_{\text{PDI}} = -42.44^\circ$, is markedly less extreme than the Case 2 optimum: the constrained search settles near the boundary of the admissible region rather than deep inside it, consistent with the expected behaviour that introducing an inequality constraint biases the optimum towards its active boundary. Second, the velocity residual of the optimal seed, $\max_k c_k = 0.14$, remains dominated by the $|v_{\text{PDI}}|$ component as in Case 2, confirming that the hemispheric restriction has not altered the structural velocity-versus- ΔV trade-off identified in Section 6.3.3: the low- ΔV cluster continues to admit arrival

Table 6.5. Multi-start convergence for Case 3 (southern-hemisphere arrival): ΔV^* and operational violation per seed.

Seed	ΔV^* [km/s]	$\max_k c_k$	t_{PSO} [min]
1	0.8994	2.0×10^{-3}	2.4
2	0.7685	1.4×10^{-1}	1.4
3	0.8404	6.3×10^{-2}	1.2
4	0.8886	$\leq 10^{-6}$ (nominal feasible)	0.6
5	0.7746	1.3×10^{-1}	1.3
6	0.8903	$\leq 10^{-6}$ (nominal feasible)	0.7
7	0.9120	$\leq 10^{-6}$ (nominal feasible)	0.7
8	0.8872	3.4×10^{-3}	0.7
9	0.8859	$\leq 10^{-6}$ (nominal feasible)	0.6
10	0.8903	1.7×10^{-3}	0.7

speeds outside the nominal band, while the nominal-velocity cluster pays an additional ~ 130 m/s in mission ΔV .

Detailed trajectory plots are omitted for this case, as the geometry is qualitatively similar to the descent profiles of Case 2; the cumulative impact of the geographic constraints is more clearly illustrated by the strict south-polar configuration, to which the discussion now turns.

6.5 Case 4: South Polar Arrival + TOF Constraint

Case 4 introduces the stringent geographic constraint that motivates the entire Scenario 2: arrival within the south-polar band defined by $\phi_{\text{PDI}} \leq -85^\circ$. The decision vector retains the twelve-dimensional three-impulse parametrisation of Equation (6.3); the constraint set of Case 3 is tightened by replacing the hemispheric residual $\phi \leq 0$ with the polar residual $\phi \leq -85^\circ$, while the altitude, velocity, surface-impact and (still inactive) TOF constraints are inherited unchanged.

Unlike Case 3, in which the hemispheric requirement was structurally satisfied by the natural Case 2 optimum, the polar constraint of Case 4 is genuinely active: the natural arrival latitude of the unconstrained descent geometry is approximately -66° (Section 6.3), some 19° north of the polar band. The trajectory geometry must therefore be modified to bias the arrival point a substantial angular distance further south, and the ΔV cost of this modification provides a quantitative measure of the cost of polar accessibility from the NRHO 9:2 reference orbit at the chosen epoch.

6.5.1 Optimal Decision Vector

The lowest- ΔV seed satisfying $\max_k c_k \leq 0.5$ converges to $\Delta V^* = 0.8598$ km/s (seed 2), with the converged decision vector reported in Table 6.6. The flight time set by the PDI gate event is $\tau_{\text{actual}} = 0.156$ d (≈ 3.75 h), an order of magnitude shorter than the 16.5 h descent of Case 2.

The short flight time is a defining feature of the south-polar solution. To bias the

Table 6.6. Converged decision vector for Case 4 (south-polar arrival).

	Departure ($k = 1$)	Mid-course ($k = 2$)	PDI braking ($k = 3$)
$\Delta v_{k,x}$ [km/s]	+0.0245	+0.0061	+0.0074
$\Delta v_{k,y}$ [km/s]	−0.6728	+0.0150	+0.1512
$\Delta v_{k,z}$ [km/s]	+0.0061	+0.0305	−0.0141
$ \Delta \mathbf{v}_k $ [km/s]	0.6732	0.0346	0.1520
τ_k [d]	0.0541	0.0174	0.0860

arrival latitude to -85° , the PSO abandons the extended-arc geometry of Cases 2 and 3, where the spacecraft completed a substantial fraction of an orbit about the Moon before reaching the gate, in favour of a direct descent in which the departure impulse propels the spacecraft along a trajectory towards the polar region. The departure burn retains its characteristic retrograde-in-plane orientation ($-Y_{\text{sel}}$ dominant at -0.6728 km/s, 99.9% of the impulse magnitude) and increases in magnitude relative to Case 2 ($0.567 \rightarrow 0.673$ km/s, $+106$ m/s); this single additional burn at the start of the trajectory accounts for the bulk of the geometric reorientation required by the polar constraint.

6.5.2 Trajectory Geometry

The converged trajectory is shown in Figure 6.5 (3D selenocentric view) and Figure 6.6 (altitude and latitude profiles versus time).

The selenocentric trajectory of Figure 6.5 confirms the qualitative shift: the descent subtends only a one and a half orbits about the Moon, rather than the two and a half orbits observed in Case 2.

6.5.3 ΔV Budget

The total mission ΔV for the optimal seed is

$$\Delta V_{\text{tot}}^{(4)} = |\Delta \mathbf{v}_{\text{dep}}^*| + |\Delta \mathbf{v}_{\text{mcc}}^*| + |\Delta \mathbf{v}_{\text{PDI}}^*| = 0.6732 + 0.0346 + 0.1520 = 0.8598 \text{ km/s.} \quad (6.5)$$

Comparing with the hemispheric-constrained optimum of Case 3 ($\Delta V^{(3)} = 0.7685$ km/s), the imposition of the polar bound introduces an additional cost of ~ 91 m/s (+12%). This is a significant outcome of the cumulative analysis: the bulk of the ΔV penalty for accessing the south-polar band is contained in this single transition.

The distribution of the budget across the three impulses also differs qualitatively from the natural-TOF case. In Case 2, the mid-course correction contributed 20% of the total ΔV and the terminal braking only 3.6%; in Case 4, this proportion is essentially inverted, with $|\Delta \mathbf{v}_{\text{mcc}}^*| = 35$ m/s (4%) and $|\Delta \mathbf{v}_{\text{PDI}}^*| = 152$ m/s (18%). The polar geometry yields an arrival speed of 1.7430 km/s, close enough to the nominal upper bound of 1.710 km/s that the braking burn becomes operationally meaningful again. The mid-course refinement, on the other hand, has little work to do once the departure impulse has correctly oriented the trajectory towards the polar band; it reduces to a small geometric correction.

A second observation concerns the PDI velocity. With $|v_{\text{PDI}}| = 1.7430$ km/s, the




Figure 6.5. Case 4 (south-polar arrival) optimised trajectory in the selenocentric inertial frame.

residual on the velocity upper bound is reduced to 33 m/s, dramatically smaller than the ~ 170 m/s violation characteristic of Cases 2 and 3. The geometric reorientation imposed by the polar constraint thus has a beneficial secondary effect: the trajectories that satisfy the polar requirement also tend to arrive at the gate with velocities much closer to the nominal band. This near-convergence of the velocity profile, achieved without any direct tightening of the velocity constraint, indicates that the polar geometry and the prescribed entry kinetics are mutually consistent in the operational sense.

A third observation, of more general scope, concerns the time of flight evolution along the cumulative-constraint sequence. The optimal τ_{actual} values converge to progressively shorter durations as geographic constraints are tightened: from 16.5 h in Case 2 (no geographic constraint) to 13.1 h in Case 3 (southern hemisphere) and to 3.75 h in the present Case 4 (south-polar arrival). This trend reflects a fundamental geometric property of the descent problem rather than a numerical artefact. In Case 2, the unconstrained arrival latitude allows the PSO to exploit extended orbital arcs in the selenocentric equatorial plane, on which the energy is dissipated gradually at low ΔV cost. In Case 4, the polar constraint displaces the target 19° further south than the natural Case 2 arrival, and no efficient extended-time geometry remains accessible to reach the polar band. The relation between geographic target accessibility and optimal flight time is therefore non-




Figure 6.6. Selenographic latitude and altitude profiles. The dashed horizontal line at -85° in the upper panel marks the polar constraint; the converged trajectory crosses this boundary near the PDI gate event.

trivial: imposing a stringent geometric constraint can compress the optimal τ by an order of magnitude.

6.5.4 Multi-start Convergence

The multi-start convergence reported in Table 6.7 exhibits a markedly different structure than the preceding cases. The ΔV spread expands considerably, from $\Delta V \in [0.86, 1.78]$ km/s, indicating that the polar-constrained problem admits several distinct basins of attraction. Two clusters are discernible: a low- ΔV cluster (seeds 2, 3, 7, 9, 10) with $\Delta V \in [0.86, 1.16]$ km/s and residuals dominated by minor velocity excursions, and a high- ΔV cluster (seeds 1, 4–6, 8) with $\Delta V \in [1.06, 1.78]$ km/s that includes strictly nominal-feasible solutions ($\max_k c_k \leq 10^{-6}$).

The wide spread is symptomatic of the increased geometric difficulty of the problem: with the polar constraint approximately at the natural edge of the accessible descent corridor, the PSO finds multiple qualitatively different transfer arcs that each terminate near $\phi = -85^\circ$, with substantially different ΔV costs. The most expensive seeds (4 and 5, $\Delta V > 1.5$ km/s) correspond to trajectories where the spacecraft executes a more elaborate geometric manoeuvre to reach the pole while strictly meeting the velocity window; the cheapest seeds (2, 3) achieve the polar arrival on more efficient corridors at the cost of moderate velocity overshoots. The gap of ~ 920 m/s between the cheapest and most expensive seed is the largest of any case in this chapter, and indicates that the multi-start strategy is essential to reliably identify the global optimum in the polar-constrained regime.

Table 6.7. Multi-start convergence for Case 4 (south-polar arrival): ΔV^* and operational violation per seed. The lowest- ΔV seed accepted under the operational tolerance is highlighted.

Seed	ΔV^* [km/s]	$\max_k c_k$	t_{PSO} [min]
1	1.0875	1.6×10^{-1}	1.6
2	0.8598	3.3×10^{-2}	1.1
3	0.9246	2.6×10^{-3}	1.1
4	1.7837	$\leq 10^{-6}$ (nominal feasible)	0.9
5	1.5890	7.9×10^{-3}	1.4
6	1.0639	2.5×10^{-4}	1.0
7	0.9236	7.1×10^{-2}	1.1
8	1.5153	$\leq 10^{-6}$ (nominal feasible)	0.9
9	1.1622	4.8×10^{-3}	1.3
10	0.9227	4.8×10^{-3}	0.9

6.6 Case 5: Free Departure Point on the NRHO

The final case of the cumulative-constraint sequence admits the departure point on the reference NRHO as a thirteenth optimisation variable α . The decision vector becomes

$$\mathbf{x}^{(5)} = \left[\Delta \mathbf{v}_{\text{dep}}^\top, \tau_1, \Delta \mathbf{v}_{\text{mcc}}^\top, \tau_2, \Delta \mathbf{v}_{\text{PDI}}^\top, \tau_3, \alpha \right]^\top \in \mathbb{R}^{13}, \quad (6.6)$$

where $\alpha \in [0, 1)$ denotes the fractional position of the departure epoch within the apolune-to-apolune segment that brackets the reference epoch t_{ref} . Both endpoints correspond to the lunar apolune; the perilune lies at $\alpha = \alpha_{\text{peri}} \approx 0.4912$ by NRHO symmetry. The apolune-anchored parametrisation eliminates the periodic redundancy of the equivalent perilune-centred offset $[-0.5, +0.5]$ and is consistent with the analogous treatment of the free-apolune case in the LEO→NRHO scenario.

The constraint set of Case 4 is inherited unchanged: altitude window, velocity window, polar latitude bound, and surface-impact safeguard. A nominal upper bound on the total time of flight, $\tau_{\text{tot}} \leq 0.6$ d, is retained operationally in the optimisation pipeline but proves structurally inactive at every converged solution discussed below. The multi-start size is increased from $N_{\text{seeds}} = 10$ to $N_{\text{seeds}} = 50$ in view of the expanded search space.

The Case 4 problem is a strict slice of the Case 5 problem at $\alpha = \alpha_{\text{peri}}$, so the inequality $\Delta V^{(5)} \leq \Delta V^{(4)}$ must hold at the true global optimum. The role of this case is to test whether the freedom to relocate the departure point along the NRHO actually yields a ΔV saving relative to the perilune-fixed Case 4 at the reference epoch of 10 March 2028.

6.6.1 Three Representative Variants

The fifty-seed multi-start, analysed under the operational tolerance $\max_k c_k \leq 0.5$, reveals a notably richer solution landscape than any preceding case. Three qualitatively distinct converged solutions are retained as representative of the principal clusters (Table 6.8):

- **Seed 46** - short-arc nominal-velocity solution with $\Delta V^* = 0.9019$ km/s and $\tau_{\text{actual}} = 0.071$ d (1.7 h). Its geometry and impulse distribution are essentially those of the

Case 4 optimum: the optimised $\alpha = 0.4913$ deviates from the reference perilune by less than 0.01 h, and the polar arrival is reached through a direct meridional descent.

- **Seed 28** - a long-arc nominal-velocity solution with $\Delta V^* = 0.9059$ km/s and $\tau_{\text{actual}} = 0.524$ d (12.6 h). The departure point is again essentially at the perilune ($\alpha = 0.4961$), but the descent corridor selected by the PSO is qualitatively different: an extended ballistic arc subtending three quarters of an orbit around the Moon.
- **Seed 5** - a velocity-relaxing solution with $\Delta V^* = 0.6781$ km/s and $|v_{\text{PDI}}| = 2.068$ km/s, exceeding the nominal upper bound by 358 m/s. This solution is presented for diagnostic purposes only; it will be shown in Section 6.6.4 to be operationally inferior to either nominal-velocity variant once a closed accounting of the mission propellant budget is adopted.

Table 6.8. Converged decision vectors for the three representative variants of Case 5.

Variant	Phase	$\Delta v_{k,x}$	$\Delta v_{k,y}$ [km/s]	$\Delta v_{k,z}$	$ \Delta \mathbf{v}_k $ [km/s]	τ_k [d]
Seed 46 (short-arc) $\alpha = 0.4913$ (peri +0.01 h)	Dep.	+0.0068	−0.6641	+0.0126	0.6642	0.0505
	Mcc	+0.0001	+0.0006	+0.0189	0.0189	0.0147
	PDI	−0.0034	+0.1871	+0.1133	0.2188	0.2000
Seed 28 (long-arc) $\alpha = 0.4961$ (peri +0.75 h)	Dep.	−0.0278	−0.2082	+0.1653	0.2673	0.4512
	Mcc	−0.0320	−0.2133	−0.3000	0.3695	0.0434
	PDI	+0.0072	−0.1153	+0.2430	0.2691	0.1879
Seed 5 (velocity-relax.) $\alpha = 0.5037$ (peri +1.9 h)	Dep.	+0.0046	−0.1327	+0.3705	0.3936	0.1370
	Mcc	+0.0069	−0.0002	+0.0822	0.0825	0.2770
	PDI	−0.0173	+0.0387	−0.1975	0.2020	0.2000

A striking common feature of all three variants is that the optimised α falls within ± 2 h of the reference perilune. The free- α freedom is therefore not exploited to displace the departure point in any of them: independent of the geometric strategy adopted (short-arc, long-arc, velocity-relaxing), the PSO autonomously selects a departure epoch essentially coincident with the natural perilune.

The internal distribution of the ΔV budget, however, differs sharply across the three. Seed 46 loads 74% of the budget onto a single near-meridional departure impulse and assigns the mid-course correction a negligible role (2%); this is the impulse distribution of the Case 4 short-arc descent. Seed 28 partitions the budget nearly uniformly across the three impulses (30/41/30), with the mid-course correction acting as the largest contributor and saturating its -0.3 km/s lower bound on the Z component. Seed 5 places 58% on a strongly out-of-plane ($+Z_{\text{sel}}$) departure burn and retains a moderate PDI braking, with the mid-course again playing a minor role. The three variants therefore correspond to three structurally distinct geometric strategies for reaching the polar band from the same departure point.

6.6.2 Trajectory Geometries

Figures 6.7–6.9 present the three converged trajectories in the selenocentric inertial frame; the corresponding altitude and latitude profiles are reported in Figures 6.10–6.12.




Figure 6.7. Seed 46 (short-arc).




Figure 6.8. Seed 28 (long-arc).




Figure 6.9. Seed 5 (velocity-relaxing).






Figure 6.10. Seed 46 profiles. **Figure 6.11.** Seed 28 profiles. **Figure 6.12.** Seed 5 profiles.

Seed 46 reproduces the Case 4 with a short-arc descent: the spacecraft is ejected from the NRHO by a $-Y_{\text{sel}}$ -dominated retrograde departure burn, follows a quasi-meridional ballistic trajectory towards the southern hemisphere, and reaches the PDI gate after only 1.7 h. The latitude profile shows a monotonic southward sweep from the perilune position ($\phi_0 \approx +90^\circ$) to $\phi_{\text{PDI}} = -84.99^\circ$.

Seed 28 unfolds along an extended ballistic arc, such as for Case 4 but with a $+0.75$ h excursion that approaches a half-revolution about the Moon before reaching the PDI gate.

Seed 5 occupies an intermediate geometric regime ($\tau_{\text{actual}} = 0.422$ d, 10.1 h) but is qualitatively distinguished by the absence of an effective braking burn: the small $-Z$ component of $\Delta \mathbf{v}_{\text{PDI}}$ (-0.198 km/s) reduces the inertial speed only modestly and the spacecraft enters the PDI gate at $|\mathbf{v}_{\text{PDI}}| = 2.068$ km/s. The polar latitude is reached as in seeds 46 and 28, but at a velocity well outside the nominal entry window.

6.6.3 ΔV Budget and Net-Energy Accounting

The transfer ΔV for the three variants is

$$\Delta V_{\text{tot}}^{(5, \text{seed } 46)} = 0.6642 + 0.0189 + 0.2188 = 0.9019 \text{ km/s}, \quad (6.7)$$

$$\Delta V_{\text{tot}}^{(5, \text{seed } 28)} = 0.2673 + 0.3695 + 0.2691 = 0.9059 \text{ km/s}, \quad (6.8)$$

$$\Delta V_{\text{tot}}^{(5, \text{seed } 5)} = 0.3936 + 0.0825 + 0.2020 = 0.6781 \text{ km/s}. \quad (6.9)$$

The two nominal-velocity variants (seeds 46 and 28) are within 4 m/s of each other in transfer ΔV , while seed 5 is ~ 224 m/s lower than either. Three observations frame the interpretation of these numbers.

Comparison with Case 4. The perilune-fixed Case 4 optimum is $\Delta V^{(4)} = 0.8598$ km/s (Section 6.5); the two nominal-velocity variants of Case 5 exceed it by approximately 42 and 46 m/s respectively. As discussed above, the inequality $\Delta V^{(5)} \leq \Delta V^{(4)}$ must hold at the true global optimum because $\Delta V^{(4)}$ is a feasible value of the Case 5 problem. The observed positive residual of $\sim 5\%$ is consistent with the multi-start convergence noise expected for a hybrid PSO+SQP algorithm in 13 dimensions; it does not, however, exclude the possibility that the true free- α optimum at this reference epoch is marginally lower than the Case 4 value and that further multi-start sampling or extended PSO settings would recover it. Nevertheless, the physical content of the result is unambiguous: the free- α freedom does not unlock a qualitatively cheaper descent corridor than the perilune-fixed Case 4 at the reference epoch of 10 March 2028.

Multimodality of the nominal cluster. The near-coincidence of seeds 46 and 28 in transfer ΔV (separation 4 m/s, well within PSO noise) but their radically different trajectory geometries (short-arc, 1.7 h, vs. long-arc, 12.6 h) is a non-trivial finding of the Case 5 analysis. The PSO has identified two qualitatively distinct geometric solutions that achieve the polar arrival at essentially identical propellant cost. These two variants are not equivalent in operational character: the short-arc solution offers minimal exposure time and tight predictive control of the PDI gate state, while the long-arc solution requires sustained navigation accuracy over a half-day descent but admits a substantially more relaxed impulse distribution (30/41/30 vs. 74/2/24). The two solutions populate distinct basins of attraction in the 13-dimensional optimisation problem; the selection between them is operational rather than energetic.

Net-energy accounting for the velocity-relaxing variant. The optimisation objective minimises the sum of impulsive ΔV during the transfer phase only, and does not account for the kinetic-energy budget delivered to the powered-descent guidance loop. When the velocity upper bound is relaxed through the operational tolerance, the spacecraft arrives at the PDI gate with surplus inertial energy that must ultimately be dissipated by the descent engines themselves. To first order, the propellant cost of this terminal deceleration is approximately equal to the velocity excess. For seed 5,

$$\Delta V_{\text{net}}^{(\text{seed } 5)} \approx \underbrace{0.678}_{\text{transfer}} + \underbrace{0.358}_{\text{powered-descent compensation}} = 1.036 \text{ km/s}, \quad (6.10)$$

making seed 5 approximately 134 m/s more expensive than either nominal variant once the total propellant budget is accounted for, despite its apparent advantage in transfer ΔV .

6.6.4 Operational Selection of the Reference Solution

The net-energy accounting of Equation (6.10) excludes seed 5 from the cluster of representative solutions: any closed accounting of the mission propellant budget penalises the 358 m/s velocity excess by an equivalent additional ΔV downstream of the PDI gate, with the net cost rising to ~ 1.04 km/s. Seeds 46 and 28 are by contrast operationally indistinguishable in ΔV (4 m/s separation, well within PSO noise), and the selection between them must rest on other criteria.

First, the substantially shorter actual flight time (1.7 h vs. 12.6 h) is operationally advantageous: it reduces ephemeris-error accumulation during the transfer, minimises the exposure window between the departure burn and the PDI gate, and admits a simpler mission timeline. Second, the impulse distribution of seed 46 places the bulk of the ΔV on the departure manoeuvre, executed at the perilune where attitude and station-keeping accuracy are at their best, and the residual on the terminal braking, executed close to the PDI gate where the velocity target is most stringent. Seed 28 is consequently retained as a variant representative of the long-arc basin of attraction, illustrating the multimodality of the free- α problem, but is not adopted as the reference solution.

The reference solution of Case 5 is therefore $\Delta V^{(5)} = \Delta V^{(5, \text{seed } 46)} = 0.9019$ km/s, with the understanding that an equivalent-cost long-arc variant exists and that the velocity-relaxing low- ΔV cluster is operationally disqualified.

6.6.5 Multi-start Convergence

The fifty-seed multi-start exhibits a substantially richer solution landscape than any preceding case. The ΔV values span the interval [0.68, 2.89] km/s, a range of 2.2 km/s, and several physically distinct clusters of solutions are discernible (Table 6.9).

Table 6.9. Multi-start convergence for Case 5 (free- α).

Seed	ΔV^* [km/s]	$\max_k c_k$	$ v_{\text{PDI}} $ [km/s] (Δv_{PDI} [m/s])	α	Cluster
5	0.6781	3.6×10^{-1}	2.068 (+358)	0.504	velocity-relaxing
14	0.7592	2.2×10^{-1}	1.933 (+223)	0.493	velocity-relaxing
18	0.8955	1.8×10^{-2}	1.728 (+18)	0.488	near-nominal short-arc
46	0.9019	6.4×10^{-3}	1.7105 (+0.5)	0.491	nominal short-arc (ref.)
28	0.9059	8.4×10^{-5}	1.7101 (+0.1)	0.496	nominal long-arc
11	1.7283	1.1×10^{-3}	1.711 (+1)	0.500	secondary nominal
6	1.3742	6.5×10^{-1}	2.358 (+648)	0.852	non-perilune local min.
50	1.8501	5.2×10^{-1}	2.227 (+517)	0.993	near-apolune local min.

The distribution of α values across the fifty seeds is informative independently of the ΔV ranking. Approximately 70% of the seeds (35 of 50) converge to $\alpha \in [0.48, 0.52]$, i.e. within ± 4 h of the reference perilune; the remaining 30% spread across the rest of the segment ($\alpha = 0.57, 0.62, 0.68, 0.85, 0.99$, etc.), with no cluster of comparable density.

This distribution constitutes empirical evidence that, at the reference epoch, the perilune is the unique strongly-attractive optimum of the free- α problem; the secondary clusters ($\alpha \neq \alpha_{\text{peri}}$) represent local minima of substantially higher cost ($\Delta V > 1.4$ km/s) that the multi-start strategy successfully identifies as suboptimal.

The persistence of the velocity-relaxing cluster across multiple seeds (5, 14, 18, etc.) confirms its structural nature: it is a recurrent attractor of the penalty-augmented objective, not an artefact of a single seed. The introduction of a more penalizing velocity-weighted criterion is expected to suppress this attractor automatically and is left to future work.

6.7 Secular study

It was deemed necessary to perform a secular-variation study, analogous to the one carried out for Scenario 1 (Section 5.9), in order to verify that the reference perilune is indeed the optimal departure point for the free- α formulation of Case 5 across the mission timeframe, and not merely an artefact of the single reference epoch analysed in Section 6.6. For each of the twenty periods, the representative solution was selected, among the operationally feasible seeds, as the one whose PDI arrival speed lies closest to the nominal value: this avoids biasing the secular comparison towards low- ΔV velocity-relaxing solutions whose hidden downstream propellant cost would make the transfer- ΔV figure alone misleading. Figure 6.13 shows the optimised departure phase α^* against the period index $p = 0, \dots, 19$, together with the full admissible range $\alpha \in [0, 1)$ of the apolune-to-apolune segment. Across the twenty periods, α^* remains tightly clustered around the tabulated reference-perilune position ($\bar{\alpha} = 0.4908$, $\sigma_{\alpha} = 0.0281$), with 80% of the periods converging within ± 0.02 of the local α_{peri} value. This confirms, over the roughly four-month window covered by the secular study, the structural finding already established for the reference epoch: the perilune remains the dominant attractor of the free- α optimisation regardless of the specific NRHO period considered. The only notable departure from this pattern occurs at $p = 5$ ($\alpha^* = 0.3915$), where the selected solution corresponds to a velocity-relaxing variant of the type discussed in Section 6.6.3, whose departure geometry deviates from the perilune-anchored optimum precisely because the net-energy-corrected ranking still admits a residual velocity excess at that period.




Figure 6.13. Optimised departure phase α^* versus NRHO period index for the Case 5 (free- α) secular-variation study.

6.8 Cumulative Summary

The five cases analysed in the preceding sections form a coherent cumulative-constraint progression in which each case retains the active constraint set of its predecessor and introduces a single new requirement or additional degree of freedom. The converged best solutions of the five cases are collected in Table 6.10 to consolidate the comparison and to make the marginal effect of each constraint or parametrisation change directly readable. All quantities reported refer to the lowest- ΔV seed of the multi-start campaign that satisfies the operational tolerance $\max_k c_k \leq 0.5$ defined at the beginning of this chapter.

Table 6.10. Cumulative summary of the five cases analysed in Chapter 6.

Case	New element	n_{vars}	ΔV^* [km/s]	$\Delta(\Delta V)$ [m/s]	τ_{act} [h]	$ v_{\text{PDI}} $ [km/s]	φ_{PDI} [deg]
1	Single-impulse baseline	4	1.0021	—	0.83	1.774	+38.6
2	Three-impulse parametr.	12	0.7419	−260	16.5	1.881	−66.4
3	$\varphi_{\text{PDI}} \leq 0^\circ$	12	0.7685	+27	~ 13	~ 1.85	−42.4
4	$\varphi_{\text{PDI}} \leq -85^\circ$	12	0.8598	+91	3.75	1.743	−85.3
5	Free- α departure (NRHO phase)	13	0.9019	+42	1.7	1.711	−85.0

Several structural trends emerge from the cumulative comparison and deserve explicit emphasis.

Three-impulse benefit. The transition from Case 1 to Case 2 confirms the operational utility of the three-impulse parametrisation: the admission of the mid-course correction

and the terminal braking impulses reduces the mission cost by 260 m/s relative to the single-impulse baseline, a saving of 26%. This benefit is the largest single ΔV reduction of the entire cumulative chain, and is recovered without introducing any operational constraint - it is purely a consequence of the enriched parametrisation. The subsequent cases progressively erode this margin through the imposition of cumulative restrictions.

Marginal cost of geographic restrictions. The southern-hemisphere constraint (Case 3) entails a small marginal cost of ~ 27 m/s (+3.6% relative to Case 2), as expected: the natural arrival latitude of the unconstrained descent geometry ($\varphi_{\text{PDI}} = -66^\circ$, Case 2) lies already within the admissible region, and the constraint is structurally inactive at the optimum. The polar restriction (Case 4), in contrast, requires the arrival latitude to be biased $\sim 19^\circ$ further south than the natural value and incurs a substantially larger marginal cost of 91 m/s (+12% relative to Case 3). The bulk of the geographic-constraint penalty for accessing the south-polar band is concentrated in this single transition.

Time-of-flight compression. The actual flight time decreases monotonically along the cumulative chain, from 16.5 h in the natural-TOF Case 2 to only 1.7 h in the free- α Case 5. This compression is not the direct consequence of the TOF constraint but reflects the geometric collapse of the descent corridor as geographic restrictions tighten: the polar arrival demands a direct meridional descent rather than an extended ballistic arc. The free- α formulation of Case 5 sharpens this regime further, identifying a short-arc solution with flight time below two hours.

PDI velocity convergence towards the nominal band. A non-trivial side effect of the cumulative restrictions is the progressive convergence of the arrival velocity towards the nominal upper bound $|v_{\text{PDI}}| = 1.710$ km/s. The unconstrained geometries of Cases 2 and 3 exceed the bound by ~ 170 and ~ 140 m/s respectively, while the polar constraint of Case 4 reduces this excess to only 33 m/s, and the free- α Case 5 converges essentially on the bound itself (+0.5 m/s). The polar geometry and the prescribed entry kinetics are therefore mutually consistent in the operational sense: the trajectories that meet the polar requirement also tend to arrive at the gate with velocities within the nominal band, without any further tightening of the velocity constraint.

Findings of the free- α case. The introduction of the NRHO departure phase α as a thirteenth optimisation variable in Case 5 does not yield a ΔV saving relative to the perilune-fixed Case 4; the converged minimum is ~ 42 m/s higher, consistent with the multi-start convergence noise of a hybrid PSO+SQP algorithm in 13 dimensions (Section 6.6.3). The physical content of the result is that, at the reference epoch of 10 March 2028, the perilune is the unique strongly-attractive optimum of the free- α problem (70% of the 50-seed multi-start converges within ± 4 h of the reference perilune), and no qualitatively cheaper descent corridor exists in the admissible NRHO segment. A secondary but operationally relevant finding is the identification of a long-arc nominal-velocity variant (Section 6.6.1) with ΔV equivalent to the short-arc reference solution but flight time of 12.6 h: the existence of this alternative illustrates the multimodality of the free- α problem and suggests that the selection between short-arc and long-arc descent geometries is a matter of operational rather than energetic preference.

Aggregate ΔV envelope. Considered as a whole, the cumulative-constraint sequence frames the operationally viable south-polar descent from the NRHO at the reference epoch in the range $\Delta V \in [0.86, 0.90]$ km/s for the geographically meaningful cases (4 and 5), with the single-impulse baseline of Case 1 (1.00 km/s) and the natural-TOF Case 2 (0.74 km/s) bracketing respectively the upper and lower limits of the broader cumulative range. These figures provide a quantitative reference for the ΔV budget of a south-polar descent from the near-rectilinear halo orbit.

7 | Conclusions

This thesis has addressed the optimisation of multi-impulse transfers in the cislunar environment through a hybrid particle-swarm and sequential-quadratic-programming framework, applied to two complementary mission segments: the trans-lunar leg from a Low Earth parking orbit to the near-rectilinear halo orbit (Scenario 1, Chapter 5), and the descent leg from the NRHO to the Powered Descent Initiation gate above the lunar south pole (Scenario 2, Chapter 6). The two scenarios were treated under a unified cumulative-constraint methodology in which each case retains the active constraint set of its predecessor and introduces a single new requirement or additional degree of freedom, so that the marginal ΔV cost of each operational restriction can be quantitatively isolated. This concluding chapter consolidates the principal findings, identifies the methodological lessons of broader applicability, and outlines the natural extensions of the present work.

7.1 Principal Results

Scenario 1: LEO $\rightarrow$ NRHO. The four-impulse parametrisation (TLI, two mid-course corrections, NRHO insertion) was analysed across a cumulative chain comprising the perigee-altitude baseline, time-of-flight, launch-site inclination, Van Allen radiation transit, and free-NRHO-arrival-point cases, for both the Cape Canaveral and Kourou launch sites. The reference Cape Canaveral solution converges to $J^* = 3.821$ km/s (Case 5a), of which 3.04 km/s is the trans-lunar injection ΔV_{circ} and ~ 0.78 km/s is the on-board budget distributed across the NRHO-insertion and the two mid-course correction nodes. The corresponding Kourou solution at the same reference epoch is $J^* = 3.842$ km/s (Case 5b), only 21 m/s above the Cape value but qualitatively distinct in its impulse architecture: the Kourou optimum collapses to an essentially two-impulse design (TLI + NRI ≈ 800 m/s) with the two mid-course corrections vanishing to a few m/s each. The free-arrival-point relaxation thus enables an architectural transition from a three-impulse cislunar architecture to a near-two-impulse design that is operationally near-optimal for Cape Canaveral and substantially advantageous for Kourou (-425 m/s relative to the apolune-fixed Case 4b, approximately 84% of the Kourou plane-change penalty).

The secular-variation study over twenty consecutive NRHO periods (Section 5.9) revealed a qualitative contrast between the two launch sites. The Cape Canaveral optimum is structurally stable across the window ($\sigma_{J^*} = 27$ m/s, $\sigma_{\alpha^*} = 0.07$) and the cost landscape admits a single dominant basin of attraction centred near apolune. The Kourou optimum, conversely, is multi-basin and substantially more dispersed ($\sigma_{J^*} = 120$ m/s, $\sigma_{\alpha^*} = 0.32$), with the optimal arrival phase scattering over almost the full orbital period. The dispersion ratio of approximately four-to-one in σ_{J^*} between the two launch sites is the natural

consequence of the inclination mismatch between the Kourou window ($i \in [4^\circ, 8^\circ]$) and the NRHO orbital plane: the constrained launch window forces a plane change that is sensitive to the precession of the NRHO across the synodic month, while the Cape Canaveral window ($i \in [26^\circ, 31^\circ]$) naturally aligns with the NRHO geometry and produces a structurally stable optimum.

Scenario 2: NRHO $\rightarrow$ PDI. The descent transfer from the NRHO perilune to the PDI gate was analysed under the three-impulse parametrisation across a cumulative chain comprising the single-impulse baseline, the three-impulse formulation, the southern-hemisphere bound, the south-polar bound, and the free-departure-point case. The reference operational solution at the south-polar arrival is $\Delta V^* = 0.86\text{--}0.90$ km/s (Cases 4 and 5), corresponding to a flight time of approximately 1.7–3.75 h between perilune departure and PDI gate triggering. The polar arrival was found to introduce the dominant marginal cost of the cumulative chain (+91 m/s relative to the southern-hemisphere case), while the southern-hemisphere bound is structurally inactive (+27 m/s, since the natural arrival latitude of the unconstrained descent is already at -66°). The free- α formulation does not yield a ΔV saving relative to the perilune-fixed Case 4 at the reference epoch: the converged minimum is ~ 42 m/s higher, consistent with the multi-start convergence noise of the hybrid algorithm in 13 dimensions, and the PSO autonomously settles on a departure phase essentially coincident with the natural perilune (70% of the 50-seed multi-start converges within ± 4 h of the reference perilune). The principal qualitative finding of the free- α case is the identification of a long-arc nominal-velocity variant with ΔV equivalent to the short-arc reference solution but flight time of 12.6 h, demonstrating that the descent problem admits multiple geometrically distinct basins of attraction at essentially identical propellant cost.

The perilune-attractor behaviour of the free- α optimum was further verified not to be an artefact of the specific reference epoch through a secular-variation study over twenty consecutive NRHO periods spanning approximately four months. The optimised departure phase remains tightly clustered around the tabulated perilune across the entire window ($\bar{\alpha} = 0.4908$, $\sigma_\alpha = 0.0281$, with 80% of periods converging within ± 0.02 of the local perilune), confirming that the reference-epoch finding generalises across the mission timeframe rather than being a one-off coincidence.

A non-trivial structural trend that emerges from the cumulative chain is the progressive convergence of the arrival velocity towards the nominal upper bound $|v_{\text{PDI}}| = 1.710$ km/s: the unconstrained Case 2 exceeds the bound by ~ 170 m/s, the polar Case 4 reduces this excess to 33 m/s, and the free- α Case 5 converges essentially on the bound itself. The polar geometry and the prescribed entry kinetics are therefore mutually consistent in the operational sense: the trajectories that meet the latitude requirement also tend to arrive at the gate with velocities within the nominal band, without any direct tightening of the velocity constraint.

7.2 Methodological Findings

Beyond the specific numerical results of the two scenarios, several methodological observations emerged from the cumulative analysis that are likely to be relevant to broader trajectory-optimisation problems in the cislunar environment.

First, the cumulative-constraint methodology proved well suited to isolating the marginal cost of individual operational requirements, provided that the constraints are introduced one at a time and the multi-start strategy is applied with deterministic seeds reused consistently across all cases. The use of identical PSO and SQP configurations throughout both scenarios (Section 4.5) ensured that the case-to-case variation of ΔV^* is attributable to the change of the active constraint set and not to solver-side artefacts.

Second, the multi-start size required for reliable identification of the global optimum depends strongly on the dimensionality of the decision space and on the multimodality of the cost landscape. Ten independent seeds were sufficient for the four- and twelve-dimensional fixed-perilune cases of both scenarios, but the free-arrival-point configurations required twenty seeds (Scenario 1, Cases 5a and 5b) or fifty seeds (Scenario 2, Case 5) to characterise the multi-basin structure of the 13-dimensional landscape.

Third, the operational feasibility tolerance is not a one-size-fits-all parameter. For Scenario 1, the strict tolerance $\|c\|_\infty \leq 10^{-3}$ proved adequate in all cases. For Scenario 2, the velocity-window constraint on $|v_{\text{PDI}}|$ generates a structural trade-off between transfer ΔV and arrival kinetic energy that is not captured by a single feasibility tolerance: tight tolerances exclude low- ΔV velocity-relaxing solutions that are operationally relevant when the downstream propellant cost is accounted for, while loose tolerances accept solutions with hidden net-energy penalties. The relaxed operational tolerance $\max_k c_k \leq 0.5$ adopted throughout Chapter 6 was a pragmatic compromise; a velocity-weighted feasibility criterion that explicitly penalises the velocity excess by an equivalent additional ΔV would be a more principled substitute and is identified as a natural refinement of the framework.

Fourth, the free-phase parametrisation itself proved numerically delicate. The original offset design in Scenario 1, centred on the apolune ($\alpha \in [-0.5, +0.5]$), has a periodic redundancy at its boundaries, and a naive `mod()` implementation of this periodicity introduced a wraparound artefact that corrupted the free- α results. The fix was to anchor $\alpha \in [0, 1)$ at the point diametrically opposite the one causing the redundancy: Scenario 1 re-anchors $\alpha = 0$ at the perilune (with the apolune falling near $\alpha = 0.5$), while Scenario 2 re-anchors $\alpha = 0$ at the apolune (with the reference perilune falling near $\alpha = 0.5$). In both cases the point of physical interest is displaced away from the segment boundary, removing the redundancy by construction, and this re-anchoring strategy is recommended as the more robust default for phase parametrisations on closed near-rectilinear orbits.

7.3 Future Work

Several extensions of the present work follow directly from the limitations of the current analysis.

The Scenario 2 framework should be supplemented with a heavier velocity-weighted criterion that penalises more arrivals outside the nominal $|v_{\text{PDI}}|$ band by an equivalent additional ΔV downstream of the gate, as discussed in Section 7.2. The introduction of this criterion is expected to suppress the velocity-relaxing low- ΔV cluster identified in Case 5 and to yield a cleaner multi-start distribution. The framework could then be extended to include the powered-descent phase itself, which has been treated here as out of scope: a self-consistent optimisation of the descent profile from perilune to lunar touchdown, with explicit modelling of finite-burn thrust and gravity-loss penalties, would provide a complete characterisation of the propellant budget for a lunar landing.

Moreover, a secular study extended to the full Artemis IV launch window (approximately one year) and encompassing both launch sites would provide a comprehensive characterisation of the launch-epoch sensitivity of the cislunar ΔV budget.

Furthermore, a natural extension on the methodological side is the integration of finite-burn thrust into the impulsive framework. The impulsive approximation adopted throughout this thesis is appropriate for chemical-propulsion missions, but the trans-lunar leg of a realistic mission may benefit from a high-specific-impulse finite-burn refinement of the optimised impulse vectors, which would quantify the gravity-loss penalty and could be performed as a post-processing step on the converged impulsive solutions without modifying the PSO+SQP pipeline.

Finally, the extension of the framework to alternative cislunar orbital regimes, such as Distant Retrograde Orbits, butterfly orbits, or Earth-Moon L_1 halo orbits, would broaden the applicability of the cumulative-constraint methodology beyond the specific NRHO of the Artemis program. The hybrid PSO+SQP framework is structurally independent of the choice of target orbit; only the SPICE ephemerides, the box bounds and the constraint definitions need to be updated for each new operational scenario.

7.4 Final Remarks

The work presented in this thesis confirms the operational viability of hybrid PSO+SQP optimisation for multi-impulse cislunar transfers and provides quantitative reference figures for the ΔV budget of an Artemis-era mission profile: approximately 3.82 km/s for the LEO $\rightarrow$ NRHO leg from Cape Canaveral with full operational constraints, plus 0.90 km/s for the NRHO $\rightarrow$ south-polar descent leg, for a total mission ΔV of approximately 4.72 km/s delivered through a four-impulse trans-lunar architecture and a three-impulse descent architecture. The cumulative-constraint methodology adopted throughout the analysis isolates the marginal cost of each operational restriction and exposes the structural trade-offs that govern the optimisation landscape: the geometric coupling between the launch site and the NRHO for the trans-lunar leg, the transfer- ΔV -arrival-velocity trade-off for the descent leg. These structural findings transcend the specific numerical results and are expected to inform the design of future cislunar mission profiles.

Bibliography

- [1] Brian Patrick. High-fidelity low-thrust trajectory optimization to the lunar gateway. Master's thesis, University of Illinois Urbana-Champaign, 2023. Advisor: Prof. Robyn Woollands.
- [2] NASA. Nasa adds mission to artemis lunar program, updates architecture. <https://www.nasa.gov/news-release/nasa-adds-mission-to-artemis-lunar-program-updates-architecture/>, February 2026. Accessed: April 2026.
- [3] Emily M. Zimovan-Spreen, Kathleen C. Howell, and Diane C. Davis. Near rectilinear halo orbits and nearby higher-period dynamical structures: orbital stability and resonance properties. *Acta Astronautica*, 172:94–110, 2020.
- [4] Todd Gardner et al. Capstone: A summary of a highly successful mission in the cislunar environment. In *37th Annual Small Satellite Conference*, 2023.
- [5] Riccardo Roberto. Hybrid optimization technique for Earth–Moon transfer trajectories: Human and cargo missions to the lunar Gateway. Master's thesis in aerospace engineering, Politecnico di Torino, 2025. Supervisors: Prof. Lorenzo Casalino, Ing. Andrea D'Ottavio, Dr. Giorgio Fasano.
- [6] Charles S. Galbreath. Securing cislunar space and the first island off the coast of earth. Policy Paper 45, Mitchell Institute for Aerospace Studies, Air & Space Forces Association, Arlington, VA, January 2024.
- [7] Diane C. Davis, Sagar A. Bhatt, Kathleen C. Howell, Jiann-Woei Jang, Ryan L. Whitley, Fred D. Clark, Davide Guzzetti, Emily M. Zimovan, and Gregg H. Barton. Orbit maintenance and navigation of human spacecraft at cislunar near rectilinear halo orbits. In *27th AAS/AIAA Space Flight Mechanics Meeting*, number AAS 17-269, San Antonio, TX, February 2017. NASA NTRS 20170001347.
- [8] NASA. The artemis plan: NASA's lunar exploration program overview. Technical report, National Aeronautics and Space Administration, September 2020.
- [9] NASA Johnson Space Center. Artemis II mission milestones: An image and video recap. NASA Official Website, April 2026. Accessed: April 29, 2026.
- [10] NASA. Why NRHO: The artemis orbit. Technical report, National Aeronautics and Space Administration, 2023. Architecture Concept Review 2022 – White Paper.

- [11] Ryan Whitley, Diane C. Davis, Laura M. Burke, Brian P. McCarthy, Roland J. Power, Melissa L. McGuire, and Kathleen C. Howell. Earth-moon near rectilinear halo and butterfly orbits for lunar surface exploration. In *AAS/AIAA Space Flight Mechanics Meeting*, 2018. AAS 18-406.
- [12] Carlos Sánchez Lara and Joan Pau Sánchez Cuartielles. Near rectilinear orbits around the moon as operational orbit for the future deep space gateway. In *ESA – 7th International Workshop on Astrodynamics Tools and Techniques*, 2018.
- [13] Emily M. Zimovan. Characteristics and design strategies for near rectilinear halo orbits within the earth-moon system. Master’s thesis, Purdue University, 2017. MSAAE Thesis.
- [14] Diane E. Lee. Gateway destination orbit model: A continuous 15 year NRHO reference trajectory. Technical report, National Aeronautics and Space Administration, August 2019. NASA White Paper, NTRS accession no. 20190030294.
- [15] Emily M. Zimovan-Spreen, Stephen T. Scheuerle, Brian P. McCarthy, Diane C. Davis, and Kathleen C. Howell. Baseline orbit generation for near rectilinear halo orbits. In *AAS/AIAA Space Flight Mechanics Meeting*, 2023. AAS 23-268.
- [16] NASA. Cislunar autonomous positioning system technology operations and navigation experiment (capstone). Technical Report NTRS 20230018489, NASA, 2023.
- [17] Bradley Cheetham et al. Capstone: A unique cubesat platform for a navigation demonstration in a near rectilinear halo orbit. In *AIAA Space and Astronautics Forum and Exposition (SPACE 2022)*, 2022.
- [18] Aaron R. Krause, Tara L. Speacht, Yu Zhang, Charles H. Lang, and Henry J. Donahue. Simulated space radiation sensitizes bone but not muscle to the catabolic effects of mechanical unloading. *PLOS ONE*, 12(8):e0182403, August 2017.
- [19] Thomas H. Mader et al. Spaceflight associated neuro-ocular syndrome (sans) and the neuro-ophthalmologic issues of prolonged microgravity exposure on iss. *npj Microgravity*, 6(7), 2020.
- [20] Yuri Martin Paez et al. Spaceflight associated neuro-ocular syndrome (sans): a systematic review. *Acta Ophthalmologica*, 99, 2020.
- [21] Adrian LeBlanc, Victor Schneider, Linda Shackelford, Sylvia West, Victor Oganov, Andrei Bakulin, and Lev Veronin. Bone mineral and lean tissue loss after long duration space flight. *Journal of Musculoskeletal and Neuronal Interactions*, 7(1):67–68, 2007.
- [22] T. A. Carstens and M. J. Campola. Total dose modeling of the transit through the van allen belts. Technical Report 20220011257, NASA Goddard Space Flight Center, 2022. Presented at the Advanced Power Systems for Deep Space Exploration Conference.
- [23] Charles H. Acton. Ancillary data services of NASA’s navigation and ancillary information facility. *Planetary and Space Science*, 44(1):65–70, 1996.

- [24] Charles H. Acton, Nathaniel Bachman, Boris Semenov, and Edward Wright. A look towards the future in the handling of space science mission geometry. *Planetary and Space Science*, 150:9–12, 2018.
- [25] NASA. Nasa’s lunar exploration program overview. Technical report, National Aeronautics and Space Administration, 2020. NASA-SP-2020-921.
- [26] W. Li and M. K. Hudson. Earth’s Van Allen radiation belts: From discovery to the Van Allen Probes era. *Journal of Geophysical Research: Space Physics*, 124:8319–8351, 2019.
- [27] Luigi Mascolo. *Low-Thrust Optimal Escape Trajectories from Lagrangian Points and Quasi-Periodic Orbits in a High-Fidelity Model*. Phd thesis, Politecnico di Torino, Torino, Italy, 2024.
- [28] Emily M. Zimovan, Kathleen C. Howell, and Diane C. Davis. Near rectilinear halo orbits and their application in cis-lunar space. In *Proceedings of the 3rd IAA Conference on Dynamics and Control of Space Systems (DyCoSS3)*, number IAA-AAS-DyCoSS3-125, Moscow, Russia, 2017.
- [29] R. Sridharan, S. M. Ahmed, T. P. Das, P. Sreelatha, P. Pradeepkumar, N. Naik, and G. Supriya. Direct evidence for water (H₂O) in the sunlit lunar ambience from CHACE on MIP of Chandrayaan I. *Planetary and Space Science*, 58(6):947–950, 2010.
- [30] Mackenzie E. Mangette and Roshan T. Eapen. Exploration of direct transfer opportunities and loitering strategies for lunar surface access. *Advances in Space Research*, 77:2165–2182, 2026.
- [31] Dennis Ghisolfo. Ottimizzazione di traiettorie di discesa sulla luna con metodi diretti. Tesi di laurea magistrale, Politecnico di Torino, Torino, Italy, 2025. Corso di Laurea Magistrale in Ingegneria Aerospaziale, Relatore: Prof. Lorenzo Casalino.
- [32] John T. Betts. *Practical Methods for Optimal Control Using Nonlinear Programming*. SIAM, Philadelphia, PA, 2nd edition, 2010.
- [33] Lorenzo Casalino. Methods for space trajectory optimization. Course material, 2025. Lecture notes, Politecnico di Torino.
- [34] Bruce A. Conway. The problem of spacecraft trajectory optimization. In Bruce A. Conway, editor, *Spacecraft Trajectory Optimization*, chapter 1, pages 1–17. Cambridge University Press, 2010.
- [35] Bruce A. Conway and Stephen W. Paris. Spacecraft trajectory optimization using direct transcription and nonlinear programming. In Bruce A. Conway, editor, *Spacecraft Trajectory Optimization*, chapter 3, pages 37–78. Cambridge University Press, 2010.
- [36] Derek F. Lawden. *Optimal Trajectories for Space Navigation*. Butterworths, London, 1963.

- [37] James Kennedy and Russell Eberhart. Particle swarm optimization. In *Proceedings of ICNN'95 – International Conference on Neural Networks*, volume 4, pages 1942–1948, 1995.
- [38] Irene Giardina. Collective behavior in animal groups: Theoretical models and empirical studies. *HFSP Journal*, 2(4):205–219, 2008.
- [39] Yuhui Shi and Russell Eberhart. A modified particle swarm optimizer. pages 69–73, 1998.
- [40] Mauro Pontani and Bruce A. Conway. Particle swarm optimization applied to impulsive orbital transfers. *Acta Astronautica*, 74:141–155, 2012.
- [41] The MathWorks, Inc. Particle swarm optimization – `particleswarm`. MATLAB Global Optimization Toolbox Documentation, 2024.
- [42] Jorge Nocedal and Stephen J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.
- [43] The MathWorks, Inc. *Optimization Toolbox User's Guide*. The MathWorks, Inc., Natick, MA, 2024. Version R2024b.
- [44] Bruce A. Conway, editor. *Spacecraft Trajectory Optimization*. Cambridge University Press, 2010.
- [45] Francesco Topputo. On optimal two-impulse earth-moon transfers in a four-body model. *Celestial Mechanics and Dynamical Astronomy*, 117(3):279–313, 2013.
- [46] Behçet Açikmeşe and Scott R. Ploen. Convex programming approach to powered descent guidance for mars landing. volume 30, pages 1353–1366, 2007.
- [47] NASA. Artemis IV. <https://www.nasa.gov/mission/artemis-iv/>, 2025. Accessed: 15 May 2026.
- [48] Gaza George, Stuart P. Space radiation measurements during the Artemis I lunar mission. *Nature*, 634(8032):48–52, 2024.