



**Politecnico
di Torino**

Politecnico di Torino

Department of Industrial Engineering
Master of Science in Aerospace Engineering

Optimal Control Approach to Probabilistic Collision Avoidance under Orbital Perturbations

Candidate:

Alessia Tocca

Supervisors:

Prof. Manuela Battipede

Prof. Luigi Mascolo

Abstract

The rapid growth of satellite constellations and the increasing congestion of low Earth orbit are intensifying the need for reliable and fuel-efficient collision avoidance strategies. When a potential conjunction is detected, spacecraft operators must rapidly determine corrective maneuvers capable of guaranteeing a safe miss distance while minimizing propellant consumption and operational impact. In this context, optimal control theory provides a powerful framework for computing efficient avoidance maneuvers under strict dynamical and operational constraints.

This work presents an optimal-control approach for the design of minimum-propellant collision avoidance maneuvers in low Earth orbit. The proposed methodology combines collision risk assessment with indirect optimal control techniques. A collision map is first constructed to analyze the relative encounter geometry and determine the radial displacement required to guarantee a safe separation at the predicted collision epoch. This displacement is then imposed as the terminal constraint of a fixed-final-time optimal control problem, in which the spacecraft must achieve the prescribed radial offset while minimizing propellant consumption.

The optimal maneuver is computed using an indirect method based on Pontryagin's Maximum Principle, which provides high numerical accuracy and enables a detailed analysis of the adjoint variables governing the optimal thrust direction. The investigation is first conducted in an ideal two-body framework to extract the reference solution for a short chemical-propulsion maneuver. The dynamical model is then progressively enriched by introducing lunisolar gravitational perturbations and Earth non-sphericity, allowing a systematic assessment of how environmental disturbances influence both the optimal trajectory and the costate structure.

Results show that, for short chemical maneuvers, the optimal thrust arc remains predominantly tangential and the velocity costates remain nearly constant even in the presence of perturbations, leading to an almost time-invariant thrust direction. The analysis is further extended to electric propulsion scenarios, including Hall-effect and ion thrusters. In these regimes the maneuver duration increases significantly, resulting in a distributed thrust profile and a more pronounced evolution of the thrust direction driven by the velocity costates.

The proposed framework provides a systematic methodology for designing fuel-efficient collision avoidance maneuvers under different propulsion architectures and dynamical environments, contributing to the development of robust guidance strategies for spacecraft operating in increasingly congested orbital regimes.

Acronyms

CA	Collision Avoidance.
EGM	Earth Gravitational Model.
EGM2008	Earth Gravitational Model 2008.
EME2000	Earth Mean Equator and Equinox of J2000.
EOM	Equations of Motion.
ESA	European Space Agency.
ICRF	International Celestial Reference Frame.
J2000	Julian epoch J2000.
LEO	Low Earth Orbit.
LS	Lunisolar perturbations.
MPBVP	Multi-Point Boundary Value Problem.
OCP	Optimal Control Problem.
OCULUS	Optimal Control with User-friendly Layout Unified Solver.
PMP	Pontryagin's Maximum Principle.
RTN	Radial-Transverse-Normal.
TCA	Time of Closest Approach.
TPBVP	Two-Point Boundary Value Problem.

ZEN

Zenith–East–North (local reference frame).

Nomenclature

Symbol	Description	Value	Unit
a	Semimajor axis	–	m
\mathbf{a}_p	Total perturbing acceleration	–	m/s ²
$(a_p)_u, (a_p)_v, (a_p)_w$	Components of \mathbf{a}_p in the local ZEN frame	–	m/s ²
\mathbf{a}_J	Earth non-sphericity perturbing acceleration	–	m/s ²
\mathbf{a}_{LS}	Lunisolar perturbing acceleration	–	m/s ²
\mathbf{a}_b	Third-body perturbing acceleration	–	m/s ²
$(a_b)_u, (a_b)_v, (a_b)_w$	Components of \mathbf{a}_b in the local ZEN frame	–	m/s ²
a_{\max}	Maximum spacecraft acceleration	–	m/s ²
b	Generic third body	–	–
c	Effective exhaust velocity	–	m/s
\mathbf{C}	Covariance matrix in the encounter plane	–	–
\bar{C}_{nm}	Fully-normalized Stokes cosine coefficients	–	–
d	Magnitude of the relative position vector $\boldsymbol{\rho}$	–	m
\mathbf{d}	Relative vector $\mathbf{r}'_{Eb} - \mathbf{r}'$ in the local ZEN frame	–	m
D^2	Squared Mahalanobis distance	–	–
\mathcal{D}	Hard-body disk in the encounter plane	–	–
e	Eccentricity	–	–

Symbol	Description	Value	Unit
e_1, e_2, e_3	Orthonormal basis vectors of the encounter-plane frame	–	–
f	State-dynamics vector field in the general optimal control formulation	–	–
F	Shooting residual vector	–	–
g_0	Standard gravitational acceleration	9.80665	m/s ²
\mathcal{H}	Hamiltonian function of the OCP	–	–
i	Inclination	–	rad
$\hat{\mathbf{i}}, \hat{\mathbf{j}}, \hat{\mathbf{k}}$	Unit vectors of the EME2000 inertial frame	–	–
$I_u^{(n)}, I_v^{(n)}, I_w^{(n)}$	Auxiliary harmonic summations	–	–
I_{sp}	Specific impulse	–	s
j	Arc index in the MPBVP formulation	–	–
J	Discrete cost function for planar target selection	–	–
\mathcal{J}	Cost functional of the OCP	–	–
\mathcal{J}^*	Augmented cost functional of the OCP	–	–
K	Number of points in the discrete time grid	–	–
k	Mahalanobis level parameter	–	–
k^*	Index of the discrete minimum of the relative distance	–	–
m	Spacecraft mass	–	kg
	Number of control variables in the general optimal control formulation	–	–
m_0	Initial spacecraft mass	462.0	kg

Symbol	Description	Value	Unit
N	Maximum degree and order of the truncated spherical-harmonic expansion	8	–
	Monte Carlo sample size	–	–
n	Degree of the spherical-harmonic expansion of the geopotential	–	–
	Dimension of the state vector in the general optimal control formulation	–	–
N_s	Number of imposed switching times in the shooting formulation	–	–
n_p	Number of arcs in the multi-point boundary value problem formulation	–	–
P_c	Probability of collision	–	–
$P_c^{(0)}$	Collision probability at the time of closest approach	–	–
$P_c^{(1)}, P_c^{(2)}$	Collision probabilities on the subsequent encounter planes	–	–
\mathbf{p}	Relative position projected onto the encounter plane	–	m
	Design vector of the shooting formulation	–	–
$\mathbf{p}^{(i)}$	i -th Monte Carlo sample projected onto the encounter plane	–	m
\mathbf{p}'	Projected relative position after a planar displacement	–	m
\bar{P}_{nm}	Fully-normalized associated Legendre function	–	–
$\bar{P}'_{nm}, \bar{P}''_{nm}$	First and second derivatives of \bar{P}_{nm}	–	–
$\hat{\mathbf{p}}_{SC}, \hat{\mathbf{q}}_{SC}, \hat{\mathbf{w}}_{SC}$	Unit vectors of the spacecraft perifocal frame	–	–
q	Number of independent boundary constraints in the general optimal control formulation	–	–
r	Geocentric distance	–	m
r_0	Initial geocentric distance	–	m
r_f	Final geocentric distance	–	m
$\mathbf{r}_s, \mathbf{r}_d$	Satellite and debris position vectors in the inertial frame	–	m

Symbol	Description	Value	Unit
\mathbf{r}	Spacecraft position vector with respect to the Earth center	–	m
\mathbf{r}_{Eb}	Earth-to-third-body position vector in the inertial frame	–	m
r_{Eb}	Magnitude of the Earth-to-body position vector \mathbf{r}_{Eb}	–	m
\mathbf{r}'_{Eb}	Components of \mathbf{r}_{Eb} in the local ZEN frame	–	m
\mathbf{r}_{bSC}	Spacecraft-to-body relative position vector in the inertial frame	–	m
r_{bSC}	Magnitude of the spacecraft-to-body relative position vector \mathbf{r}_{bSC}	–	m
\mathbf{r}'_{bSC}	Components of \mathbf{r}_{bSC} in the local ZEN frame	–	m
R	Equivalent hard-body radius	20	m
R_E	Earth reference equatorial radius	6378136.3	m
$\mathbf{R}(\vartheta, \varphi)$	Rotation matrix from the EME2000 inertial frame to the local ZEN frame	–	–
\bar{S}_{nm}	Fully-normalized Stokes sine coefficients	–	–
$\mathcal{S}_{\mathcal{F}}$	Switching function associated with thrust and coast arcs	–	–
t	Time	–	s
t_{ref}	Reference epoch	January 1, 2000 12:00:00 UTC	UTC
t_0	Initial time of the propagation grid	–	s
t_k	k -th node of the discrete time grid	–	s
t_{TCA}	Time of closest approach	–	s
$t_{\text{TCA,ref}}$	Refined time of closest approach	–	s
t_1, t_2	Subsequent times for robustness assessment	–	s
	Switching times	–	s
$t_{(j-1)+}, t_{j-}$	Initial and final time values of the j -th arc in the MPBVP formulation	–	s
t_f	Final time of the OCP	–	s
t_{start}	Optimized maneuver start time	–	s
t_{coll}	Collision time imposed as terminal epoch	–	s

Symbol	Description	Value	Unit
t_{man}	Natural maneuver duration from the free-final-time formulation	–	s
\mathbf{T}	Thrust vector	–	N
T	Magnitude of the thrust vector \mathbf{T}	–	N
T_{max}	Maximum thrust magnitude	–	N
$\hat{\mathbf{T}}$	Unit vector of the thrust direction	–	–
T_u, T_v, T_w	Components of \mathbf{T} in the local ZEN frame	–	N
T_{nm}, U_{nm}	Harmonic longitude blocks	–	–
\mathbf{u}	Control vector in the general optimal control formulation	–	–
\mathcal{U}	Admissible control set in the general optimal control formulation	–	–
u, v, w	Velocity components in the local ZEN frame	–	m/s
$\hat{u}, \hat{v}, \hat{w}$	Unit vectors of the local ZEN frame	–	–
V	Total gravitational potential of the Earth gravity model	–	m^2/s^2
\mathbf{V}	Orthogonal matrix whose columns are the eigenvectors of the covariance matrix \mathbf{C}	–	–
	Kinematic vector appearing in the compact Hamiltonian expression	–	–
$\mathbf{v}_s, \mathbf{v}_d$	Satellite and debris velocity vectors in the inertial frame	–	m/s
$\mathbf{v}_1, \mathbf{v}_2$	Eigenvectors of the covariance matrix \mathbf{C} in the encounter plane	–	–
\mathbf{v}_{rel}	Relative velocity vector defined as $\mathbf{v}_d - \mathbf{v}_s$	–	m/s
x	Auxiliary variable defined as $x = \sin \varphi$	–	–
\mathbf{x}	Spacecraft state vector in spherical ZEN coordinates	–	–
	State vector in the general optimal control formulation	–	–
\mathbf{x}_0	Initial state vector for the Monte Carlo propagation	–	–
	Initial state vector of the optimal control problem	–	–

Symbol	Description	Value	Unit
\mathbf{x}_f	Final state vector of the OCP	–	–
$\mathbf{x}_{(j-1)+}, \mathbf{x}_{j-}$	Initial and final state vectors of the j -th arc in the MPBVP formulation	–	–
y_k	Sampled squared relative distance, $y_k = d_k^2$	–	m ²
$\widehat{y}(t)$	Quadratic interpolation of the squared relative distance	–	m ²
\mathbf{y}	Augmented state-costate vector in the boundary-value problem formulation	–	–
α_T	Elevation angle of the thrust direction in the local ZEN frame	–	rad
β_T	Heading angle of the thrust direction in the local ZEN frame	–	rad
δ_p	Deterministic planar displacement in the encounter plane	–	m
δ_p^*	Selected planar displacement target	–	m
Δm	Propellant mass consumed during the maneuver	–	kg
Δt_{burn}	Burn duration of the active thrust arc	–	s
Δt_j	Duration of the j -th arc in the MPBVP formulation	–	s
Δv	Velocity increment generated by the maneuver	–	m/s
$\Delta \vartheta$	Terminal difference in geocentric longitude between optimized and nominal trajectories	–	rad
$\Delta \varphi$	Terminal difference in geocentric latitude between optimized and nominal trajectories	–	rad
$\bar{\varepsilon}$	Mean absolute error in the Monte Carlo validation of the collision probability	–	–
ε_{max}	Maximum absolute error in the Monte Carlo validation of the collision probability	–	–

Symbol	Description	Value	Unit
\mathcal{E}	Specific mechanical energy	–	m^2/s^2
θ	Angular variable in the polar parameterization of the hard-body disk or covariance ellipse	–	rad
ϑ	Geocentric longitude	–	rad
ϑ_{lo}	Terrestrial longitude	–	rad
$\vartheta_{\text{G,ref}}$	Greenwich sidereal angle at the reference epoch t_{ref}	280.46061837504	deg
λ	Adjoint vector in the optimal control formulation	–	–
$\lambda_r, \lambda_\vartheta, \lambda_\varphi$	Adjoint variables associated with the position coordinates	–	–
$\lambda_u, \lambda_v, \lambda_w$	Adjoint variables associated with the velocity components	–	–
λ_m	Adjoint variable associated with the spacecraft mass	–	–
λ_V	Adjoint velocity vector used in the switching function and thrust-direction law	–	–
λ_V	Magnitude of the adjoint velocity vector λ_V	–	–
$\lambda_{(j-1)+}, \lambda_{j-}$	Initial and final adjoint vectors of the j -th arc in the MPBVP formulation	–	–
λ_1, λ_2	Eigenvalues of the covariance matrix in the encounter plane	–	m^2
μ	Generic gravitational parameter	–	m^3/s^2
μ_E	Earth gravitational parameter	$3.98600441500 \times 10^{14}$	m^3/s^2
μ_ℓ	Moon gravitational parameter	$4.90280107600 \times 10^{12}$	m^3/s^2
μ_s	Sun gravitational parameter	$1.32712440018 \times 10^{20}$	m^3/s^2

Symbol	Description	Value	Unit
$\boldsymbol{\mu}$	Mean vector of the relative position projected onto the encounter plane	–	m
	Vector of Lagrange multipliers associated with the boundary constraints	–	–
ν	True anomaly	–	rad
ρ	Auxiliary ratio defined as $\rho = R_E/r$	–	–
$\boldsymbol{\rho}$	Relative position vector defined as $\boldsymbol{r}_d - \boldsymbol{r}_s$	–	m
	Relative position vector evaluated at the time of closest approach	–	m
ρ_d	Magnitude of the vector \boldsymbol{d}	–	m
ρ_b	Magnitude of the Earth-to-body position vector \boldsymbol{r}_{Eb}	–	m
$\Sigma_{x_0}^s$	Initial covariance matrix of the spacecraft state	–	–
$\Sigma_{x_0}^d$	Initial covariance matrix of the debris state	–	–
$\boldsymbol{\sigma}_s$	Vector of 1σ uncertainties of the spacecraft initial state	–	–
$\boldsymbol{\sigma}_d$	Vector of 1σ uncertainties of the debris initial state	–	–
σ_r	Standard deviation of the initial radial coordinate	–	m
$\sigma_\vartheta, \sigma_\varphi$	Standard deviations of the initial angular coordinates	–	rad
$\sigma_u, \sigma_v, \sigma_w$	Standard deviations of the initial velocity components	–	m/s
σ_m	Standard deviation of the initial mass	–	kg
φ	Geocentric latitude	–	rad
φ	Mayer term of the OCP	–	–
Φ	Disturbing potential defined as $\Phi = V - \mu_E/r$	–	m^2/s^2
Φ	Lagrange term of the OCP	–	–

Symbol	Description	Value	Unit
ω	Argument of periapsis	–	rad
ω_E	Earth rotation rate used in the definition of the terrestrial longitude	$7.29211514671 \times 10^{-5}$	rad/s
ξ_1, ξ_2	Principal coordinates in the covariance-eigenvector basis	–	m
ξ	Projected relative position in the encounter plane expressed in principal coordinates	–	m
Ω	Right ascension of the ascending node	–	rad
$\Pi(\cdot)$	Projection operator onto the encounter plane	–	–
\mathcal{X}	Constraint vector of the OCP	–	–

Contents

List of Figures	16
List of Tables	18
1 Introduction	19
1.1 Preface	19
1.2 The Space Debris Problem	20
1.3 Optimization Methods in Space Trajectory Design	23
1.4 Objectives and Contributions	24
1.5 Structure of the Thesis	25
2 Dynamical Model and Perturbing Accelerations	27
2.1 Reference Frames and State Definition	28
2.1.1 EME2000 and Local ZEN Frame	28
2.1.2 State Vector Definition	30
2.2 Dynamical Model in Spherical ZEN Coordinates	30
2.3 Perturbing Accelerations	32
2.3.1 Earth Non-Sphericity Model (EGM2008)	32
2.3.1.1 Geopotential Expansion	32
2.3.1.2 Fully-Normalized Legendre Functions	34
2.3.1.3 Recurrence Relations	34
2.3.1.4 Derivatives for the Costate Equations	35
2.3.2 Lunisolar Perturbations	37
2.3.2.1 Third-Body Formulation	37
2.3.2.2 Projection in ZEN Coordinates	38
2.3.2.3 Derivatives for the Costate Equations	39
2.4 Complete Perturbed Equations of Motion	40
3 Conjunction Geometry and Probability of Collision	42
3.1 Short-Encounter Assumption	43
3.2 Relative Motion and Time of Closest Approach	44
3.2.1 Discrete Search	44
3.2.2 Parabolic Refinement	45
3.3 Encounter Plane Definition	46
3.3.1 Plane Perpendicular to Relative Velocity	47
3.4 Gaussian Model in the Encounter Plane	48
3.4.1 Projection of State and Covariance	48
3.4.2 Mean and Covariance Estimation	50
3.4.3 Covariance Ellipse and Mahalanobis Distance	51
3.4.4 Integration over the Hard-Body Disk	54

3.5	Probability Maps and Planar Displacement	55
3.5.1	Planar Displacement at the Time of Closest Approach	55
3.5.2	Numerical Quadrature and Construction of the Probability Map	56
3.5.3	Subsequent Encounter Planes and Robustness	59
3.5.4	Target Displacement Selection	60
3.6	Monte Carlo Estimate of $P_c(\delta_p)$ (Validation)	62
4	Optimal Control Problem Formulation	64
4.1	General structure of the optimal control problem	65
4.1.1	Merit index and augmented formulation	66
4.1.2	Boundary conditions for optimality	67
4.1.3	Equations for adjoint and control variables	69
4.1.4	Multi-Point Boundary Value Problem Formulation	70
4.2	Spacecraft Trajectory Optimal Control Formulation	73
4.2.1	Adopted Dynamical Model and Problem Definition	73
4.2.2	Hamiltonian and Switching Function	74
4.2.3	Optimal Thrust Direction and Angular Parametrization	75
4.2.4	Switching Structure and Imposed Arc Strategy	77
4.2.5	Adjoint Equations	77
4.2.6	Numerical Solution via Shooting and Differential Correction	79
5	Optimal Solution in the Unperturbed Dynamical Model	81
5.1	Implementation within the OCULUS Solver	82
5.2	Propulsion Model and Parameter Selection	83
5.2.1	Baseline Propulsion Configuration	83
5.2.2	Continuation toward Electric Propulsion Regimes	84
5.3	Baseline Circular-Orbit Configuration	85
5.3.1	Free Final Time Formulation: Problem Definition and Switching Structure	86
5.3.2	Converged Solution	88
5.4	Continuation to the Inclined-Orbit Case	90
5.4.1	Free Final Time Formulation	91
5.4.2	Converged Solution	92
5.5	Imposition of the Selected Planar Displacement	93
5.5.1	Free Final Time Formulation	94
5.5.2	Converged Solution	94
5.6	Alignment with the Collision Time	96
5.6.1	Fixed Final Time Reformulation and Optimization of the Manuever Start	96
5.6.2	Converged Fixed-Time Solution	98
5.6.3	Angular Consistency Verification	100
5.6.4	Thrust Direction and Velocity Costates Behavior	101
5.7	Electric-Propulsion Cases	103
5.7.1	Hall-Effect Thruster Case: $T_{\max} = 0.05$ N, $c = 20000$ m/s	104
5.7.2	Ion Thruster Case: $T_{\max} = 0.01$ N, $c = 30000$ m/s	106
5.7.3	Comparative Discussion	108
6	Influence of Perturbations on the Optimal Solution	110
6.1	Lunisolar Perturbations Case	111

6.1.1	Backward Propagation	111
6.1.2	Converged Solution	112
6.1.3	Thrust Direction and Velocity Costates Behavior with Lunisolar Perturbations	114
6.2	Lunisolar Perturbations and Earth Non-Sphericity Case	115
6.2.1	Backward Propagation	115
6.2.2	Intermediate Three-arc Solution	116
6.2.3	Final Two-arc Solution	117
6.2.4	Thrust Direction and Velocity Costates Behavior with Lunisolar Perturbations and Earth Non-Sphericity	119
6.3	Comparative Discussion	120
6.3.1	Loss of Orbit Equivalence and Role of Absolute Epoch	120
6.3.2	Non-circularity and Appearance of $u \neq 0$	122
6.3.3	Increased Sensitivity and Convergence Implications	123
6.4	Electric-Propulsion Cases with Lunisolar Perturbations and Earth Non-Sphericity	123
6.4.1	Hall-Effect Thruster Case: $T_{\max} = 0.05$ N, $c = 20000$ m/s	125
6.4.2	Ion Thruster Case: $T_{\max} = 0.01$ N, $c = 30000$ m/s	126
6.4.3	Comparative Discussion	128
7	Conclusions	130
7.1	Interpretation of Results	130
7.2	Model Limitations	133
7.3	Possible Extensions	134

List of Figures

1.1	Number of tracked objects in Earth orbit over time.	21
1.2	Spatial density of space objects in LEO as a function of altitude for different size ranges.	22
2.1	EME2000 $\{\hat{\mathbf{I}}, \hat{\mathbf{J}}, \hat{\mathbf{K}}\}$, perifocal $\{\hat{\mathbf{p}}_{SC}, \hat{\mathbf{q}}_{SC}, \hat{\mathbf{w}}_{SC}\}$, and ZEN $\{\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{w}}\}$ reference systems.	28
2.2	Spacecraft position, velocity, and thrust components resolved in the local ZEN frame.	31
3.1	Relative distance profile and parabolic refinement of the time of closest approach (TCA).	46
3.2	Monte Carlo cloud of the projected relative position at the TCA.	50
3.3	Covariance ellipses of the projected relative position distribution.	52
3.4	Collision probability map $P_c(\delta_p)$ over the displacement grid.	57
3.5	Zoomed view of the probability map in the region of interest.	58
3.6	Collision probability map with contour levels highlighting regions of decreasing P_c as the planar displacement varies.	58
3.7	Projected encounter geometry on the subsequent planes used to assess temporal robustness.	60
3.7a	Encounter plane at $t_{TCA} + 60$ s.	60
3.7b	Encounter plane at $t_{TCA} + 120$ s.	60
4.1	Thrust orientation in the ZEN reference frame and definition of elevation angle α_T and heading angle β_T	76
5.1	Thrust direction and velocity costate behavior during the active thrust phase.	101
5.1a	Thrust direction angles in the RTN frame and thrust–velocity angle during the burn arc.	101
5.1b	Velocity costates $\lambda_u, \lambda_v, \lambda_w$ and their norm during the burn arc.	101
5.2	Thrust direction and velocity costate behavior during the active thrust phase. Electric-propulsion case with $T_{max} = 0.05$ N and $c = 20000$ m/s.	105
5.2a	Thrust-direction angles in RTN and thrust–velocity angle during the burn.	105
5.2b	Velocity costates and their norm during the burn.	105
5.3	Thrust direction and velocity costate behavior during the active thrust phase. Electric-propulsion case with $T_{max} = 0.01$ N and $c = 30000$ m/s.	107
5.3a	Thrust-direction angles in RTN and thrust–velocity angle during the burn.	107
5.3b	Velocity costates and their norm during the burn.	107

6.1	Thrust direction and velocity costate behavior during the active thrust phase. Lunisolar perturbations case.	114
6.1a	Thrust-direction angles in the RTN frame and thrust-velocity angle during the burn.	114
6.1b	Velocity costates and their norm during the burn.	114
6.2	Thrust direction and velocity costate behavior during the active thrust phase. Lunisolar perturbations and Earth non-sphericity case.	119
6.2a	Thrust-direction angles in the RTN frame and thrust-velocity angle during the burn.	119
6.2b	Velocity costates and their norm during the burn.	119
6.3	Thrust direction and velocity costate behavior during the active thrust phase. Electric-propulsion case with $T_{\max} = 0.05$ N and $c = 20000$ m/s under lunisolar perturbations and Earth non-sphericity.	126
6.3a	Thrust-direction angles in RTN and thrust-velocity angle during the burn.	126
6.3b	Velocity costates and their norm during the burn.	126
6.4	Thrust direction and velocity costate behavior during the active thrust phase. Electric-propulsion case with $T_{\max} = 0.01$ N and $c = 30000$ m/s under lunisolar perturbations and Earth non-sphericity.	128
6.4a	Thrust-direction angles in RTN and thrust-velocity angle during the burn.	128
6.4b	Velocity costates and their norm during the burn.	128

List of Tables

- 3.1 Nominal initial states of satellite and debris. 48
- 3.2 1σ uncertainties assigned to the initial state components. 49

- 5.1 Circular equatorial case – free-final-time baseline optimal solution. . . 90
- 5.2 Inclined circular case ($i = 30^\circ$) – free-final-time baseline optimal solution. 92
- 5.3 Inclined circular case ($i = 30^\circ$) – free-final-time optimal solution for a 105 m radial displacement. 95
- 5.4 Inclined circular case ($i = 30^\circ$) – fixed-final-time optimal solution aligned with the collision epoch. 99
- 5.5 Inclined circular case ($i = 30^\circ$) – fixed-final-time solution with Hall-effect propulsion. 104
- 5.6 Inclined circular case ($i = 30^\circ$) – fixed-final-time solution with ion propulsion. 106

- 6.1 Inclined circular case ($i = 30^\circ$) – fixed-final-time solution including lunisolar perturbations. 112
- 6.2 Inclined circular case ($i = 30^\circ$) – intermediate three-arc solution used for convergence regularization, including lunisolar perturbations and Earth non-sphericity. 117
- 6.3 Inclined circular case ($i = 30^\circ$) – final fixed-final-time solution including lunisolar perturbations and Earth non-sphericity 118
- 6.4 Inclined circular case ($i = 30^\circ$) – fixed-final-time solution with Hall-effect propulsion including perturbations. 125
- 6.5 Inclined circular case ($i = 30^\circ$) – fixed-final-time solution with ion propulsion including perturbations. 127

1 Introduction

1.1 Preface

Over the past decades, the near-Earth space environment has undergone a profound transformation. What was once a relatively sparse operational domain has evolved into a densely populated and dynamically complex region. The increasing reliance on satellite-based services for communication, navigation, Earth observation and security has made the sustainability of space activities a matter of strategic relevance.

The progressive increase in orbital congestion has fundamentally changed the way conjunction events are perceived and managed. In the early decades of space activity, close approaches were relatively rare and often handled on a case-by-case basis. Today, the operational landscape is characterized by frequent conjunction warnings, complex coordination between independent satellite operators, and increasingly stringent requirements on mission continuity and safety. As a consequence, collision avoidance must be treated not merely as a reactive maneuver, but as a rigorously defined engineering problem involving uncertainty, performance trade-offs, and dynamical constraints.

Within this context, collision avoidance is no longer an occasional operational adjustment, but an integral component of mission design and long-term orbital management. The growing number of active satellites, combined with legacy debris and fragmentation events, requires systematic approaches capable of integrating dynamical modeling, uncertainty assessment and trajectory optimization into a coherent framework. At the core of conjunction assessment lies the recognition that orbital states are never known with absolute precision. Tracking measurements, atmospheric density modeling, sensor limitations and force model uncertainties introduce errors that propagate over time. At the Time of Closest Approach (TCA), the relative position between two objects is therefore best described statistically rather than deterministically.

This probabilistic characterization, typically expressed through covariance matrices and collision probability estimates in the encounter plane, provides a quantitative measure of risk. However, translating such probabilistic information into a concrete maneuver strategy remains a nontrivial step that requires a consistent dynamical formulation.

This thesis addresses the problem of probabilistic collision avoidance by coupling conjunction risk assessment with the formulation and solution of an OCP. Rather than limiting the analysis to the identification of potentially hazardous encounters, the work focuses on translating probabilistic information in the encounter plane into a dynamically consistent target displacement, and subsequently into an optimal corrective trajectory. The approach adopted in this work seeks precisely to bridge the gap

between risk quantification and maneuver design. Instead of treating risk assessment and maneuver generation as separate processes, the probabilistic description of the conjunction event is embedded directly into the trajectory optimization problem.

Collision probability maps in the encounter plane are used as quantitative indicators of risk, from which a suitable displacement is selected. The displacement selected in the encounter plane is not an abstract geometric shift, but becomes a physically meaningful terminal condition within the spacecraft dynamics. The selected condition is then imposed as a terminal constraint in an indirect OCP solved via differential correction. From a methodological standpoint, the use of an indirect optimal control formulation allows the necessary conditions of optimality to be derived analytically through Pontryagin's Maximum Principle. The resulting boundary value problem is solved via differential correction techniques, ensuring that both the dynamical equations and the imposed terminal displacement are satisfied simultaneously. This strategy provides a structured way to determine thrust direction, the sequencing and duration of thrust and coast arcs, and the corresponding propellant consumption in a dynamically consistent manner.

The analysis is conducted under both unperturbed and perturbed dynamical models expressed in geocentric spherical coordinates, allowing a direct assessment of how orbital perturbations influence the resulting solution. While the two-body approximation provides a simplified reference scenario in which the structure of the optimal solution can be clearly interpreted, real orbital motion in Low Earth Orbit is affected by perturbations such as Earth's oblateness and lunisolar gravitational effects. These perturbations introduce time dependence and additional coupling between state and costate dynamics, which in turn affect the sensitivity and convergence properties of the OCP. By comparing solutions obtained under unperturbed and perturbed dynamics, the thesis highlights how model fidelity impacts both the qualitative structure of the maneuver and the numerical behavior of the indirect method.

The objective of this work is not to propose a universal operational framework, nor to address the full spectrum of space traffic management challenges. Rather, it aims to provide a coherent and technically rigorous link between probabilistic conjunction analysis and optimal trajectory design. By embedding collision probability information into a constrained OCP, the study contributes to a systematic methodology for collision avoidance in congested orbital environments.

In an era where orbital sustainability is increasingly recognized as a prerequisite for continued space utilization, the development of integrated approaches that combine uncertainty modeling and trajectory optimization represents an essential step toward more robust and analytically grounded collision avoidance strategies.

1.2 The Space Debris Problem

Since the beginning of the space age in 1957, thousands of spacecraft have been launched into Earth orbit. Alongside operational missions, a significant population of non-functional objects has accumulated, including defunct satellites, upper stages, fragmentation debris from explosions and collisions, and mission-related fragments.

Over time, the orbital environment has progressively transitioned from a relatively sparse operational domain to a dynamically complex and increasingly congested region.

According to the ESA Space Environment Report 2025 [1], more than 30,000 objects larger than 10 cm are currently tracked in Earth orbit, while the number of smaller debris fragments is estimated to be orders of magnitude higher. Objects in the centimeter and millimeter-size range are generally not catalogued systematically, yet they remain capable of inflicting severe damage due to the high relative velocities typical of orbital motion. In Low Earth Orbit (LEO), encounter velocities can reach several kilometers per second, implying that even millimetric fragments carry kinetic energies comparable to that of high-velocity projectiles. Consequently, the debris problem cannot be assessed solely in terms of object count, but must also consider the dynamical and energetic characteristics of orbital encounters.

The long-term evolution of the debris environment is strongly influenced by fragmentation events and the progressive deployment of large constellations in LEO. Historical breakups, whether accidental explosions or collisions, have generated clouds of fragments that remain in orbit for extended periods, depending on altitude and atmospheric drag conditions. At higher LEO altitudes, where atmospheric drag is weak, orbital lifetimes can span decades, leading to persistent contamination of specific altitude shells. The same ESA report highlights that, despite mitigation measures and post-mission disposal guidelines, the overall debris population continues to grow. This sustained increase raises concerns about the long-term sustainability of heavily populated orbital regions.

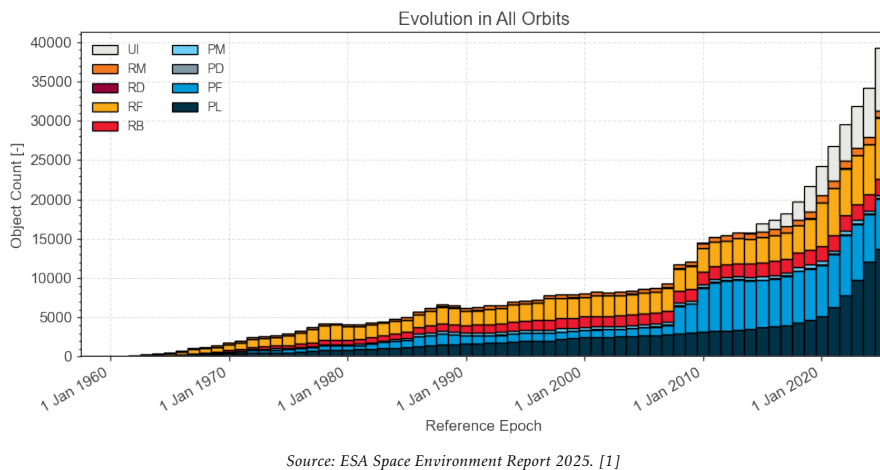



Figure 1.1: Number of tracked objects in Earth orbit over time.

Beyond the temporal growth of the population, its spatial distribution is equally relevant. The debris environment is not uniformly distributed across altitudes, but exhibits pronounced peaks in specific LEO bands. These peaks correspond to operational regimes heavily exploited for Earth observation, telecommunications constellations, and scientific missions. In such regions, active satellites coexist with legacy debris and fragments from past breakup events, resulting in locally elevated collision probabilities.

Figure 1.2 illustrates how the spatial density of objects varies with altitude. The

presence of density peaks implies that collision risk is strongly altitude-dependent, and that operational strategies must account for the local dynamical environment rather than relying on global object statistics. The coexistence of tracked objects (typically larger than 10 cm) and a significantly larger population of smaller, partially untracked fragments further complicates the assessment of conjunction risk.



Source: ESA Space Environment Report 2025. [1]

Figure 1.2: *Spatial density of space objects in LEO as a function of altitude for different size ranges.*

The debris environment also exhibits complex dynamical behavior. Relative motion between objects is governed by orbital mechanics, perturbations due to Earth’s oblateness, lunisolar gravitational effects, and atmospheric drag. Over time, these perturbations modify orbital elements, causing orbital planes and nodes to precess and altering relative geometries. As a result, conjunction conditions are inherently time-dependent and sensitive to the absolute epoch. Even small uncertainties in state estimation may lead to significant variations in predicted miss distances at the Time of Closest Approach (TCA).

The presence of uncertainty in orbit determination further complicates conjunction analysis. State estimation errors, atmospheric drag modeling inaccuracies, sensor limitations, and incomplete knowledge of small debris populations imply that close approaches cannot be assessed deterministically. Instead, they must be treated within a probabilistic framework, where the relative position at TCA is modeled as a random variable characterized by a covariance matrix. Collision avoidance therefore becomes a decision-making problem under uncertainty, where risk reduction must be balanced against propellant consumption, mission constraints, and operational costs.

From an operational perspective, the increasing congestion of LEO has led to a growing number of conjunction warnings issued to satellite operators. Only a fraction of these events require maneuver execution, yet each alert demands assessment, analysis, and often coordination between independent operators. The scalability of current operational practices is therefore challenged by the projected growth of constellations and the continued accumulation of debris.

In this context, collision avoidance can no longer be treated as an isolated corrective action, but must be integrated within a systematic framework that links probabilistic risk evaluation with dynamically consistent maneuver design. The development of such a framework constitutes the focus of the present work, where probabilistic conjunction

analysis in the encounter plane is directly connected to the formulation and solution of an OCP for collision avoidance.

1.3 Optimization Methods in Space Trajectory Design

Optimization techniques have played a central role in astrodynamics since the early development of space mission analysis. Early trajectory design approaches were predominantly based on analytical solutions derived under simplified dynamical assumptions, such as two-body motion and impulsive maneuvers. While these formulations provided valuable insight and computational efficiency, their applicability becomes limited when complex constraints, perturbations, and realistic mission requirements must be considered.

The systematic use of optimal control theory has progressively expanded the scope of trajectory design beyond purely analytical solutions. Modern mission analysis increasingly relies on numerical optimization techniques capable of handling high-fidelity dynamical models, bounded thrust profiles, nonlinear perturbations, and operational constraints [2–4]. Within this framework, OCPs are formulated in continuous time and solved either through indirect methods, based on Pontryagin’s Maximum Principle [2], or through direct numerical approaches that discretize the control and state variables [3].

Indirect methods provide the necessary conditions of optimality in the form of a Hamiltonian boundary value problem, coupling state and adjoint dynamics. These formulations allow the analytical structure of the optimal control law to be characterized explicitly and provide insight into switching behavior and arc sequencing [2, 5, 6]. Direct approaches, on the other hand, transform the continuous OCP into a finite-dimensional nonlinear programming problem and are often preferred for their numerical robustness and flexibility in the treatment of constraints [3]. Both families of methods have been successfully applied in fuel-optimal, time-optimal, and multi-phase trajectory design problems in space applications [4, 5, 7–9].

When realistic perturbations are included in the dynamical model, such as Earth’s non-sphericity, atmospheric drag, or third-body gravitational effects, the structure of the optimal solution may become significantly more complex. Nonlinear coupling between states and controls, path constraints, and control bounds can lead to discontinuous or piecewise-defined thrust profiles. As a consequence, trajectory optimization in high-fidelity orbital environments requires careful numerical treatment and consistent modeling assumptions.

In the specific context of collision avoidance, the optimization problem acquires an additional level of complexity. The maneuver must not only satisfy dynamical feasibility and propulsion limitations, but also ensure that the resulting trajectory reduces the probability of collision below an acceptable operational threshold. This requirement introduces probabilistic considerations into the control problem, effectively coupling deterministic orbital dynamics with uncertainty propagation and risk assessment models.

Unlike classical transfer optimization, where terminal conditions are typically expressed in terms of position and velocity matching, collision avoidance involves terminal requirements defined in the encounter geometry. The displacement imposed in the encounter plane must be sufficient to achieve the desired risk reduction, while minimizing propellant expenditure and limiting long-term deviations from the reference orbit. The control problem therefore becomes constrained by both dynamical and probabilistic conditions, which must be treated consistently within a unified formulation.

Traditional operational strategies often rely on heuristic in-track maneuvers or simplified geometric displacements designed to increase miss distance at the time of closest approach. Although effective in many operational scenarios, such approaches do not necessarily guarantee optimality in terms of propellant consumption, long-term trajectory deviation, or robustness under uncertainty. Embedding the avoidance requirement within a formal optimal control framework enables a systematic determination of thrust direction, timing, and duration of thrust and coast phases, while ensuring consistency with the underlying spacecraft dynamics.

By combining collision probability maps with the solution of a fuel-OCP, a displacement selected on the basis of risk metrics can be realized through a trajectory that satisfies the imposed terminal conditions while minimizing propellant consumption. In this way, probabilistic conjunction analysis and optimal control are treated within a coherent dynamical framework, extending classical trajectory optimization techniques to the specific requirements of collision avoidance in congested orbital environments.

1.4 Objectives and Contributions

The primary objective of this thesis is to develop an integrated framework for probabilistic collision avoidance that consistently connects conjunction risk assessment with optimal control under both simplified and perturbed orbital dynamics.

Rather than treating conjunction analysis and maneuver design as independent operational steps, the work aims to establish a structured link between the probabilistic characterization of a close approach and its dynamical realization through an optimal trajectory. The focus is therefore not limited to identifying hazardous events, but extends to embedding the selected risk-reduction strategy directly into the spacecraft dynamics.

More specifically, the thesis aims to:

- formulate the conjunction event in terms of relative motion and statistical uncertainty at the Time of Closest Approach (TCA);
- construct collision-probability maps in the encounter plane as quantitative risk indicators, enabling a direct visualization of how displacement in different directions affects the probability of collision;
- identify a dynamically meaningful target displacement capable of reducing collision risk according to prescribed criteria, translating probabilistic information into a terminal condition suitable for trajectory design;

- formulate and solve the resulting avoidance maneuver as an indirect OCP, in which the selected displacement is imposed as a terminal constraint and the spacecraft dynamics are expressed in geocentric spherical coordinates;
- implement the solution procedure within the OCULUS (Optimal Control with User-friendly Layout Unified Solver) framework [10], employing differential correction techniques to solve the associated boundary value problem;
- assess the influence of dynamical model fidelity by comparing solutions obtained under unperturbed two-body dynamics and under perturbed models including Earth’s non-sphericity and lunisolar gravitational effects;
- quantify how orbital perturbations modify the optimal control structure, burn timing, thrust direction evolution, and propellant consumption, and evaluate the sensitivity of the solution to the adopted dynamical assumptions.

The contribution of this work lies in the systematic integration of probabilistic conjunction analysis with indirect optimal control within a unified dynamical framework. Collision probability information, typically confined to risk assessment procedures, is here embedded directly into the trajectory optimization problem through a geometrically consistent terminal displacement in the encounter plane.

Furthermore, by analyzing both unperturbed and perturbed dynamical models expressed in spherical coordinates, the study provides insight into how model fidelity influences not only quantitative performance metrics, but also the qualitative structure of the optimal maneuver. This dual-level analysis clarifies the role of perturbations in shaping the control law and highlights the importance of consistent modeling when collision avoidance is formulated as an OCP.

In this way, the thesis contributes a coherent methodological framework for probabilistic collision avoidance that preserves dynamical consistency, explicitly accounts for uncertainty in encounter geometry, and leverages the analytical structure of indirect optimal control methods for trajectory design in congested orbital environments.

1.5 Structure of the Thesis

The remainder of this thesis is organized as follows.

Chapter 2 introduces the dynamical framework adopted throughout the work. The geocentric spherical ZEN state representation is first defined, together with the corresponding kinematic relations. The two-body baseline model is then presented as a reference configuration, followed by the introduction of selected orbital perturbations, including Earth’s non-sphericity and lunisolar gravitational effects. These elements lead to the formulation of the complete perturbed equations of motion that constitute the dynamical foundation for both conjunction analysis and optimal control.

Chapter 3 addresses the conjunction analysis problem. The encounter geometry is defined in terms of relative motion at the Time of Closest Approach (TCA), and the statistical characterization of uncertainty is introduced. The probabilistic formulation

used to compute collision probability in the encounter plane is presented, together with the procedure adopted to construct collision-probability maps. Based on these maps, a target displacement is selected according to prescribed risk-reduction criteria, translating probabilistic information into a geometrically consistent condition to be enforced within the subsequent optimal control formulation.

Chapter 4 formulates the OCP within the adopted dynamical framework. The performance index, boundary conditions, and constraints are introduced, and the indirect formulation based on Pontryagin's Maximum Principle is derived. The resulting set of necessary conditions is cast into a boundary value problem, establishing the mathematical structure used to compute the avoidance maneuver consistent with the imposed terminal displacement.

Chapter 5 presents the construction of the optimal solution in the unperturbed dynamical model. The numerical strategy implemented within the OCULUS (Optimal Control with User-friendly Layout Unified Solver) framework [10] is described, together with the continuation approach adopted to obtain convergence toward the final collision-avoidance trajectory. The final solution is then analyzed in detail, examining how burn duration, propellant consumption, and thrust direction evolve as the propulsion parameters vary.

Chapter 6 investigates the influence of orbital perturbations on the optimal solution. By repeating the analysis under the perturbed dynamical model, a direct comparison between unperturbed and perturbed cases is performed. The effects of perturbations are examined in terms of trajectory deformation, adjoint dynamics, thrust direction evolution, burn timing, and propellant consumption. Furthermore, the dependence of the solution on the propulsion parameters is reassessed under perturbed dynamics, highlighting how variations in maximum thrust and exhaust velocity modify burn duration, mass consumption, and thrust orientation when gravitational perturbations are included.

Finally, Chapter 7 concludes the thesis by providing an overall interpretation of the results. The main findings are summarized, the limitations of the adopted modeling assumptions are discussed, and possible extensions of the present work are outlined, including further refinement of dynamical models and alternative solution strategies.

2 Dynamical Model and Perturbing Accelerations

This chapter introduces the dynamical framework adopted to model the spacecraft motion and to support the subsequent conjunction analysis and optimal control formulation. The description of the spacecraft dynamics is a fundamental prerequisite for both risk evaluation and maneuver synthesis, as it provides the mathematical representation of how the state evolves under gravitational and non-gravitational influences as well as under the action of control inputs.

The goal of this chapter is to define a state representation that is well suited to thrust-guided orbital motion and indirect optimal-control methods, and to progressively enrich the baseline dynamics with the perturbing effects that are relevant in Low Earth Orbit (LEO). The resulting model combines a spherical coordinate description of the spacecraft state with local projections of forces and accelerations, enabling a unified treatment of both the nominal motion and perturbative contributions.

To this end, the chapter begins by defining the reference frames used throughout the thesis, focusing on both the inertial geocentric frame and a local moving frame attached to the spacecraft. Within this setting, the spherical ZEN state vector is introduced and interpreted in terms of physically meaningful velocity components along the local radial, eastward, and northward directions. The corresponding unperturbed equations of motion are then derived by projecting Newton's second law onto the local basis, including the contribution of thrust acceleration and the associated mass depletion law.

Once the baseline two-body model has been established, the chapter describes the perturbing accelerations considered in this work. Earth non-sphericity is modeled through a truncated spherical harmonic expansion based on the EGM2008 gravity model, while third-body perturbations due to the Sun and the Moon are included via a standard differential acceleration formulation. For both contributions, the expressions required to evaluate the perturbing accelerations in spherical ZEN variables are presented, together with the analytical derivatives needed later to construct the full state-costate system in the OCP.

The chapter concludes by collecting all terms into a single set of complete perturbed equations of motion, which constitute the dynamical model used throughout the remainder of the thesis. This comprehensive model serves as the foundation for analytical developments and numerical implementations in the subsequent chapters, ensuring consistency between the dynamical description and the optimal control framework.

2.1 Reference Frames and State Definition

2.1.1 EME2000 and Local ZEN Frame

A rigorous definition of the adopted reference systems is required before introducing the dynamical model. Although the spacecraft state will be expressed in spherical ZEN variables, an inertial geocentric frame is necessary as a global reference for the definition of angular coordinates, orbital elements, and frame transformations.

In this work, the adopted inertial reference frame is the Earth Mean Equator and Equinox of J2000 (EME2000). The EME2000 frame is centered at the Earth's center of mass; its fundamental plane coincides with the mean equatorial plane at epoch J2000; the $\hat{\mathbf{I}}$ axis is aligned with the vernal equinox direction; the $\hat{\mathbf{K}}$ axis is normal to the equatorial plane and oriented toward the North celestial pole; and the $\hat{\mathbf{J}}$ axis completes the right-handed orthonormal triad.

Under the assumptions adopted in this study, EME2000 is treated as inertial over the time scales of interest. Although EME2000 does not exactly coincide with the International Celestial Reference Frame (ICRF), the rotational offset between the two frames is on the order of a few hundredths of arcsecond [9, 11]. Such a discrepancy is negligible with respect to the propagation accuracy required in the present work, and the frame can therefore be regarded as quasi-inertial for orbital dynamics purposes.

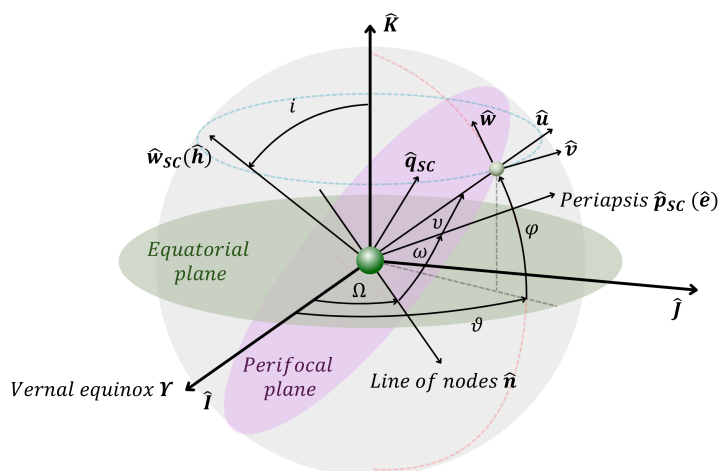


Figure 2.1: EME2000 $\{\hat{\mathbf{I}}, \hat{\mathbf{J}}, \hat{\mathbf{K}}\}$, perifocal $\{\hat{\mathbf{p}}_{sc}, \hat{\mathbf{q}}_{sc}, \hat{\mathbf{w}}_{sc}\}$, and ZEN $\{\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{w}}\}$ reference systems.

While EME2000 provides the global inertial reference, the spacecraft equations of motion are formulated in a local moving frame attached to the spacecraft, denoted ZEN (Zenith–East–North). This frame is centered at the spacecraft position and defined through geocentric spherical coordinates. Its unit vectors are constructed as follows:

- the Zenith direction $\hat{\mathbf{u}}$, aligned with the outward radial direction from the Earth's center to the spacecraft;
- the East direction $\hat{\mathbf{v}}$, tangent to the parallel of latitude and oriented toward increasing longitude;

- the North direction $\hat{\mathbf{w}}$, tangent to the meridian and oriented toward increasing geocentric latitude.

The resulting triad is orthonormal and right-handed, satisfying $\hat{\mathbf{u}} \times \hat{\mathbf{v}} = \hat{\mathbf{w}}$. In this formulation, both the spacecraft velocity and all acceleration contributions (central gravity, thrust, and perturbations) are projected onto the local radial, eastward, and northward directions. This representation is particularly convenient for the spherical dynamical model adopted in [9] and used throughout this thesis.

For completeness, transformations between the local ZEN frame and the inertial EME2000 frame are conveniently performed through an intermediate perifocal frame. The perifocal frame $\{\hat{\mathbf{p}}_{SC}, \hat{\mathbf{q}}_{SC}, \hat{\mathbf{w}}_{SC}\}$ is attached to the instantaneous Keplerian orbit of the spacecraft and is defined as follows: the unit vector $\hat{\mathbf{p}}_{SC}$ is directed from the central body toward the periapsis; the unit vector $\hat{\mathbf{w}}_{SC}$ is aligned with the specific angular momentum vector and is therefore normal to the orbital plane; the unit vector $\hat{\mathbf{q}}_{SC}$ completes the right-handed orthonormal triad and lies in the orbital plane, orthogonal to $\hat{\mathbf{p}}_{SC}$.

The orientation of the perifocal frame with respect to EME2000 is uniquely determined by the classical orbital elements. The line of nodes is defined as the intersection between the orbital plane and the Earth's equatorial plane; its direction, from the Earth's center toward the ascending node, serves as the reference for the definition of the right ascension of the ascending node Ω . This angle is measured in the equatorial plane from the inertial $\hat{\mathbf{I}}$ axis to the ascending node direction and specifies the rotation of the orbital plane about the inertial $\hat{\mathbf{K}}$ axis.

The inclination i defines the tilt of the orbital plane with respect to the equatorial plane and is measured as the angle between the orbital angular momentum vector and the inertial $\hat{\mathbf{K}}$ axis. The argument of periapsis ω determines the orientation of the periapsis within the orbital plane and is measured from the ascending node to the periapsis direction along the direction of motion.

Together with these three orientation angles, the semi-major axis a and the eccentricity e characterize the size and shape of the Keplerian orbit. The semi-major axis a represents half of the longest diameter of the orbital ellipse and directly determines the orbital energy in the two-body approximation. The eccentricity e quantifies the deviation of the orbit from circularity, with $e = 0$ corresponding to a circular orbit and $0 < e < 1$ to an elliptical trajectory.

Finally, the true anomaly ν specifies the instantaneous position of the spacecraft along the orbital path. It is defined as the angle, measured in the orbital plane, between the periapsis direction and the current position vector of the spacecraft. The set $\{a, e, i, \Omega, \omega, \nu\}$ therefore provides a complete parametrization of the Keplerian motion in the two-body approximation.

In the present work, orbital elements are used for state initialization, interpretation of the reference orbit, and consistency checks between different representations of the motion. The perifocal frame therefore plays a key intermediary role in linking the inertial description in EME2000 to the local ZEN formulation adopted for the dynamical

equations.

2.1.2 State Vector Definition

The spacecraft position is expressed through geocentric spherical coordinates referred to the EME2000 frame. Let r denote the geocentric distance, ϑ the geocentric longitude measured in the equatorial plane from the inertial $\hat{\mathbf{I}}$ axis, and φ the geocentric latitude measured from the equatorial plane toward the North direction. The domain of the angular coordinates is $\vartheta \in (-\pi, \pi]$ and $\varphi \in [-\frac{\pi}{2}, \frac{\pi}{2}]$.

The spacecraft state vector in spherical ZEN coordinates is defined as:

$$\mathbf{x} = [r \ \vartheta \ \varphi \ u \ v \ w \ m]^T,$$

where u , v , and w denote the velocity components resolved along the local Zenith, East, and North directions, respectively, and m is the spacecraft mass.

By construction, these components are related to the time derivatives of the spherical coordinates through the kinematic relations:

$$u = \dot{r}, \quad v = r \cos \varphi \dot{\vartheta}, \quad w = r \dot{\varphi}.$$

This representation allows the velocity vector to be directly decomposed into geometrically meaningful contributions: radial motion along $\hat{\mathbf{u}}$, longitudinal motion along $\hat{\mathbf{v}}$, and latitudinal motion along $\hat{\mathbf{w}}$. The adopted state formulation therefore preserves a clear physical interpretation while remaining fully consistent with the inertial reference geometry.

2.2 Dynamical Model in Spherical ZEN Coordinates

Once the state variables have been introduced, the corresponding kinematic and dynamical relations can be established.

The kinematic equations follow directly from the definition of the velocity components in the ZEN frame. Since the velocity vector is decomposed along the local radial, East, and North directions, the evolution of the spherical coordinates is governed by:

$$\dot{r} = u, \quad \dot{\vartheta} = \frac{v}{r \cos \varphi}, \quad \dot{\varphi} = \frac{w}{r}.$$

These relations are purely geometric and depend only on the spherical parametrization of position; they are therefore independent of the specific force model.

The dynamical equations are obtained by projecting Newton's second law onto the local ZEN basis. Under the assumption of a central gravitational field generated by the Earth and including a thrust acceleration term, the equations of motion take the following




Figure 2.2: Spacecraft position, velocity, and thrust components resolved in the local ZEN frame.

first-order form:

$$\begin{cases} \dot{r} = u, \\ \dot{\vartheta} = \frac{v}{r \cos \varphi}, \\ \dot{\varphi} = \frac{w}{r}, \\ \dot{u} = -\frac{\mu_E}{r^2} + \frac{v^2 + w^2}{r} + \frac{T_u}{m}, \\ \dot{v} = -\frac{uv}{r} + \frac{vw}{r} \tan \varphi + \frac{T_v}{m}, \\ \dot{w} = -\frac{uw}{r} - \frac{v^2}{r} \tan \varphi + \frac{T_w}{m}, \\ \dot{m} = -\frac{T}{c} \end{cases}$$

where μ_E denotes the Earth's gravitational parameter, T is the thrust magnitude, and c is the effective exhaust velocity. The quantities T_u , T_v , and T_w represent the projections of the thrust vector onto the local ZEN directions.

The thrust vector \mathbf{T} is decomposed in the ZEN frame as:

$$T_u = T \sin \alpha_T, \quad T_v = T \cos \alpha_T \cos \beta_T, \quad T_w = T \cos \alpha_T \sin \beta_T,$$

where α_T defines the elevation angle of the thrust vector with respect to the local horizontal plane, and β_T defines its azimuth within that plane. A detailed discussion of the role of these angles within the optimal control formulation will be provided in the corresponding section.

The formulation reported above follows the spherical ZEN representation adopted in [9]. In the present subsection, only the central two-body gravitational model has been considered. Additional perturbative contributions will be introduced in the following sections to construct the complete perturbed dynamical model.

2.3 Perturbing Accelerations

In addition to the central two-body gravitational attraction, the dynamical model adopted in this thesis accounts for selected perturbing effects that are relevant for Low Earth Orbit applications. In particular, the perturbations considered are those arising from the Earth's non-spherical gravity field, modeled through the EGM2008 spherical-harmonic expansion, and the third-body gravitational attraction of the Sun and the Moon.

The cumulative perturbing acceleration is expressed as the superposition of the individual contributions:

$$\mathbf{a}_p = \mathbf{a}_J + \mathbf{a}_{LS},$$

where \mathbf{a}_J represents the acceleration due to Earth's non-sphericity and \mathbf{a}_{LS} denotes the combined lunisolar third-body perturbation.

These additional accelerations modify the purely central gravitational dynamics introduced in the previous section and introduce time-dependent and non-central contributions to the motion. Within the spherical ZEN formulation, the perturbing acceleration vector is incorporated into the equations of motion through its projections onto the local unit vectors $\{\hat{\mathbf{u}}, \hat{\mathbf{v}}, \hat{\mathbf{w}}\}$, yielding the components $(a_p)_u$, $(a_p)_v$, and $(a_p)_w$. The following subsections present the adopted models and discuss their implementation in the ZEN dynamical framework.

2.3.1 Earth Non-Sphericity Model (EGM2008)

2.3.1.1 Geopotential Expansion

The Earth's gravity field departs from the ideal central field due to the planet's oblateness and higher-order mass distribution. A standard way to represent such effects is to introduce the gravitational potential as a spherical harmonic expansion. In this work, the Earth Gravitational Model EGM2008 is adopted in the Tide-Free system [12, 13]. The dimensional gravitational potential is written as [9, 13–16]:

$$V(r, \varphi, \vartheta) = \frac{\mu_E}{r} \left[1 + \sum_{n=2}^N \left(\frac{R_E}{r} \right)^n \sum_{m=0}^n \bar{P}_{nm}(\sin \varphi) \left(\bar{C}_{nm} \cos(m\vartheta_{lo}) + \bar{S}_{nm} \sin(m\vartheta_{lo}) \right) \right], \quad (2.1)$$

where μ_E is the Earth's gravitational parameter, R_E is the reference equatorial radius used by the model, ϑ and φ are the geocentric longitude and latitude, and \bar{C}_{nm} , \bar{S}_{nm} are the fully-normalized Stokes coefficients provided by EGM2008. Harmonics up to degree and order $N = 8$ are implemented, together with the associated fully-normalized Legendre functions \bar{P}_{nm} and the spherical harmonic coefficients \bar{C}_{nm} and \bar{S}_{nm} .

The truncation at degree and order 8 represents a compromise between model fidelity and computational cost. In low Earth orbit, the dominant contribution arises from the J_2 term associated with the Earth's oblateness, while higher-degree harmonics decay rapidly with $(R_E/r)^n$ and produce progressively smaller corrections. Increasing the truncation degree would therefore lead to a noticeable increase in computational burden due to the larger number of harmonic terms and recursive evaluations, with only

marginal impact on the dynamical evolution for the propagation scenarios considered in this work.

In the EGM2008 expansion, the longitude entering the harmonic terms is the *terrestrial* (Earth-fixed) longitude, here denoted by ϑ_{10} . This differs from the inertial longitude ϑ due to the Earth's rotation. Assuming a constant Earth rotation rate and neglecting precession and nutation effects over the propagation interval, the relationship is written as [9]:

$$\vartheta_{10}(t) = \vartheta(t) - [\vartheta_{G,\text{ref}} + \omega_E(t - t_{\text{ref}})],$$

where ω_E is the Earth's rotation rate and $\vartheta_{G,\text{ref}}$ is the Greenwich sidereal angle at the reference epoch t_{ref} (typically J2000).

It is convenient to define the disturbing (non-central) potential:

$$\Phi(r, \varphi, \vartheta) \triangleq V(r, \varphi, \vartheta) - \frac{\mu_E}{r},$$

so that the central term is treated analytically within the baseline ZEN equations, while the contribution of Φ enters as a perturbation.

The perturbing acceleration associated with the Earth's non-sphericity is obtained as the gradient of the disturbing potential Φ . In spherical coordinates and consistently with the ZEN formulation adopted in this work, the components of the non-sphericity perturbation are therefore written as [9]:

$$(a_J)_u = \frac{\partial \Phi}{\partial r}, \quad (a_J)_v = \frac{1}{r \cos \varphi} \frac{\partial \Phi}{\partial \vartheta}, \quad (a_J)_w = \frac{1}{r} \frac{\partial \Phi}{\partial \varphi}.$$

By substituting the harmonic expansion (2.1) into the definition of Φ and performing the partial derivatives term by term, one obtains the explicit expressions of the radial, longitudinal and latitudinal perturbing accelerations in terms of the fully-normalized spherical harmonics. These terms represent the non-central gravitational contributions that are added to the baseline two-body acceleration within the ZEN equations of motion.

For compactness and numerical efficiency, it is convenient to introduce the auxiliary quantities:

$$x = \sin \varphi, \quad \rho = \frac{R_E}{r}, \quad (2.2)$$

together with the harmonic longitude blocks:

$$\begin{aligned} T_{nm}(\vartheta_{10}) &= \bar{C}_{nm} \cos(m\vartheta_{10}) + \bar{S}_{nm} \sin(m\vartheta_{10}), \\ U_{nm}(\vartheta_{10}) &= -\bar{C}_{nm} \sin(m\vartheta_{10}) + \bar{S}_{nm} \cos(m\vartheta_{10}). \end{aligned}$$

The following intermediate summations are then defined:

$$\begin{aligned} I_u^{(n)} &= \sum_{m=0}^n T_{nm} \bar{P}_{nm}(x), \\ I_v^{(n)} &= \sum_{m=0}^n m U_{nm} \bar{P}_{nm}(x), \\ I_w^{(n)} &= \sum_{m=0}^n T_{nm} \cos \varphi \bar{P}'_{nm}(x), \end{aligned}$$

where $\bar{P}'_{nm}(x)$ denotes the derivative of the fully-normalized associated Legendre function with respect to its argument x .

By carrying out explicitly the derivatives of Φ with respect to r , ϑ and φ , the perturbing acceleration components associated with the Earth's non-spherical gravity field can be written as:

$$\begin{aligned} (a_J)_u &= \frac{\mu_E}{r^2} \sum_{n=2}^N \rho^n (n+1) I_u^{(n)}, \\ (a_J)_v &= \frac{\mu_E}{r^2 \cos \varphi} \sum_{n=2}^N \rho^n I_v^{(n)}, \\ (a_J)_w &= \frac{\mu_E}{r^2} \sum_{n=2}^N \rho^n I_w^{(n)}. \end{aligned} \tag{2.3}$$

2.3.1.2 Fully-Normalized Legendre Functions

The associated Legendre functions appearing in (2.1) must be consistent with the normalization adopted for the EGM2008 spherical harmonic coefficients [13]. Since the EGM2008 model provides fully-normalized Stokes coefficients, the implementation employs the corresponding fully-normalized associated Legendre functions.

Let x denote the argument of the Legendre functions, with $x \in [-1, 1]$; in the present geocentric latitude formulation, $x = \sin \varphi$. The fully-normalized functions are defined as:

$$\bar{P}_{nm}(x) = N_{nm} P_n^m(x), \quad N_{nm} = \sqrt{k \frac{(2n+1)(n-m)!}{(n+m)!}}, \tag{2.4}$$

where P_n^m are the (unnormalized) associated Legendre functions and $k = 1$ for $m = 0$ and $k = 2$ for $m > 0$. The adopted normalization scheme is consistent with the formulation described in [17].

2.3.1.3 Recurrence Relations

To evaluate $\bar{P}_{nm}(x)$ efficiently and robustly, the recurrence scheme proposed by Holmes and Featherstone [17] is employed. The computation proceeds by increasing the degree n and, for each n , the order m , using the sectorial terms as initial seeds.

Seed values. The base and first sectorial term are defined as:

$$\bar{P}_{00}(x) = 1, \quad \bar{P}_{11}(x) = \sqrt{3} \sqrt{1-x^2}.$$

For $m > 1$, the diagonal (sectorial) terms are obtained recursively as:

$$\bar{P}_{mm}(x) = \sqrt{\frac{2m+1}{2m}} \sqrt{1-x^2} \bar{P}_{m-1,m-1}(x), \quad m > 1.$$

Sub-diagonal. The first sub-diagonal ($n = m + 1$) is computed through:

$$\bar{P}_{m+1,m}(x) = \sqrt{2m+3} x \bar{P}_{mm}(x).$$

General recurrence. For $n \geq m + 2$, the general two-term recurrence reads:

$$\bar{P}_{nm}(x) = A_{nm} x \bar{P}_{n-1,m}(x) - B_{nm} \bar{P}_{n-2,m}(x), \quad n \geq m + 2,$$

with coefficients:

$$A_{nm} = \sqrt{\frac{(2n-1)(2n+1)}{(n-m)(n+m)}}, \quad B_{nm} = \sqrt{\frac{(2n+1)(n+m-1)(n-m-1)}{(2n-3)(n-m)(n+m)}}.$$

2.3.1.4 Derivatives for the Costate Equations

Although the full costate dynamics will be derived in the optimal control chapter, it is convenient to collect here the analytical derivatives of the EGM2008 perturbing accelerations, required to build the complete state-costate system.

Starting from the explicit harmonic expressions of the perturbing accelerations (2.3) and recalling the definitions (2.2), the derivatives can be obtained by direct term-wise differentiation of the spherical harmonic expansion.

Radial acceleration derivatives. The derivative with respect to the radius acts only on the explicit radial dependence of the factor $(\mu_E/r^2)\rho^n$, which produces an additional $(n+2)$ multiplier after differentiation. The resulting expressions are:

$$\frac{\partial(a_J)_u}{\partial r} = -\frac{\mu_E}{r^3} \sum_{n=2}^N \rho^n (n+1)(n+2) I_u^{(n)},$$

$$\frac{\partial(a_J)_u}{\partial \vartheta} = \frac{\mu_E}{r^2} \sum_{n=2}^N \rho^n (n+1) I_v^{(n)},$$

$$\frac{\partial(a_J)_u}{\partial \varphi} = \frac{\mu_E}{r^2} \sum_{n=2}^N \rho^n (n+1) I_w^{(n)}.$$

The longitude derivative originates from the trigonometric dependence of the harmonic blocks, while the latitude derivative follows from the chain rule through $x = \sin \varphi$.

Longitudinal acceleration derivatives. The longitudinal component contains the additional geometrical factor $1/(r \cos \varphi)$, whose explicit dependence on r and φ produces additional terms in the corresponding partial derivatives. One obtains:

$$\begin{aligned}\frac{\partial(a_J)_v}{\partial r} &= -\frac{\mu_E}{r^3 \cos \varphi} \sum_{n=2}^N \rho^n (n+2) I_v^{(n)}, \\ \frac{\partial(a_J)_v}{\partial \vartheta} &= -\frac{\mu_E}{r^2 \cos \varphi} \sum_{n=2}^N \rho^n \sum_{m=0}^n m^2 T_{nm} \bar{P}_{nm}(x), \\ \frac{\partial(a_J)_v}{\partial \varphi} &= \frac{\mu_E}{r^2} \left[\tan \varphi \frac{1}{\cos \varphi} \sum_{n=2}^N \rho^n I_v^{(n)} + \sum_{n=2}^N \rho^n \sum_{m=0}^n m U_{nm} \bar{P}'_{nm}(x) \right].\end{aligned}$$

The first term in the latitude derivative arises from the differentiation of $\sec \varphi$, whereas the second term is associated with the derivative of the fully-normalized Legendre functions.

Latitudinal acceleration derivatives. The derivatives of the latitudinal component involve both first and second derivatives of the fully-normalized associated Legendre functions:

$$\begin{aligned}\frac{\partial(a_J)_w}{\partial r} &= -\frac{\mu_E}{r^3} \sum_{n=2}^N \rho^n (n+2) I_w^{(n)}, \\ \frac{\partial(a_J)_w}{\partial \vartheta} &= \frac{\mu_E}{r^2} \sum_{n=2}^N \rho^n \sum_{m=0}^n m U_{nm} (\cos \varphi \bar{P}'_{nm}(x)), \\ \frac{\partial(a_J)_w}{\partial \varphi} &= \frac{\mu_E}{r^2} \sum_{n=2}^N \rho^n \sum_{m=0}^n T_{nm} \left[-\sin \varphi \bar{P}'_{nm}(x) + \cos^2 \varphi \bar{P}''_{nm}(x) \right].\end{aligned}$$

The appearance of $\bar{P}''_{nm}(x)$ results from the differentiation of $\cos \varphi \bar{P}'_{nm}(x)$ with respect to φ .

Derivatives of the fully-normalized associated Legendre functions. The latitude derivatives of the perturbing accelerations require the first and second derivatives of the fully-normalized associated Legendre functions $\bar{P}_{nm}(x)$, with $x = \sin \varphi$.

For the unnormalized functions $P_n^m(x)$, the classical identity reported by [18] reads:

$$(1-x^2) \frac{dP_n^m(x)}{dx} = -n x P_n^m(x) + (n+m) P_{n-1}^m(x).$$

Using the fully-normalized definition (2.4), the above identity can be rewritten in normalized form. In practical implementation, the Legendre functions themselves are evaluated through the stable recursive scheme of [17], and the derivative is expressed in a form that reuses the same recursively computed quantities. The first derivative with respect to the argument x can thus be written in numerically stable form as:

$$\bar{P}'_{nm}(x) = \frac{n x \bar{P}_{nm}(x) - \alpha_{nm} \bar{P}_{n-1,m}(x)}{x^2 - 1} \quad (2.5)$$

with:

$$\alpha_{nm} = \sqrt{\frac{(2n+1)(n+m)(n-m)}{2n-1}}.$$

The derivative with respect to latitude follows directly from the chain rule:

$$\frac{\partial \bar{P}_{nm}(\sin \varphi)}{\partial \varphi} = \cos \varphi \bar{P}'_{nm}(x), \quad x = \sin \varphi.$$

For completeness, the second derivative required in the expression of $\partial(a_J)_w/\partial\varphi$ is obtained by differentiating (2.5), yielding:

$$\bar{P}''_{nm}(x) = \frac{(n \bar{P}_{nm} + n x \bar{P}'_{nm} - \alpha_{nm} \bar{P}'_{n-1,m})(x^2 - 1) - (n x \bar{P}_{nm} - \alpha_{nm} \bar{P}_{n-1,m})2x}{(x^2 - 1)^2}$$

where $\bar{P}'_{n-1,m}$ is computed using the same normalized formula with $n \mapsto n - 1$.

The above formulation is fully consistent with the normalization adopted for the EGM2008 coefficients and, combined with the Holmes–Featherstone recurrence scheme for \bar{P}_{nm} , ensures accurate and numerically robust evaluation up to the selected truncation degree $N = 8$.

It is finally noted that the non-sphericity perturbation depends exclusively on the geometric variables (r, ϑ, φ) ; therefore, no partial derivatives with respect to (u, v, w, m) appear in the costate equations from this contribution.

2.3.2 Lunisolar Perturbations

2.3.2.1 Third-Body Formulation

The gravitational attraction of third bodies, primarily the Sun and the Moon, produces secular and long-period variations that are relevant for precise propagation and for long time horizons. In the adopted formulation, the third-body perturbation acting on the spacecraft is written as the difference between the acceleration of the spacecraft and that of the Earth due to the same perturbing body, thus removing the uniform acceleration of the Earth-centered frame [9].

Let b denote a generic perturbing body (Sun s or Moon ℓ), with gravitational parameter μ_b . Let \mathbf{r} be the spacecraft position w.r.t. Earth, and \mathbf{r}_{Eb} the position of body b w.r.t. Earth. The spacecraft-to-body vector is:

$$\mathbf{r}_{bSC} = \mathbf{r} - \mathbf{r}_{Eb}.$$

The third-body perturbing acceleration in a geocentric inertial frame is then [9]:

$$\mathbf{a}_b = -\mu_b \left(\frac{\mathbf{r}_{bSC}}{\|\mathbf{r}_{bSC}\|^3} + \frac{\mathbf{r}_{Eb}}{\|\mathbf{r}_{Eb}\|^3} \right) = \mathbf{a}_{SC,b} - \mathbf{a}_{E,b}. \quad (2.6)$$

The combined lunisolar perturbation is obtained as:

$$\mathbf{a}_{LS} = \sum_{b \in \{s, \ell\}} \mathbf{a}_b.$$

2.3.2.2 Projection in ZEN Coordinates

For integration in spherical ZEN variables, the perturbing acceleration (2.6) must be expressed in the local $\{\hat{u}, \hat{v}, \hat{w}\}$ basis. In the adopted spherical ZEN formulation, the local orthonormal basis is defined as:

$$\hat{u} = [\cos \varphi \cos \vartheta, \cos \varphi \sin \vartheta, \sin \varphi],$$

$$\hat{v} = [-\sin \vartheta, \cos \vartheta, 0],$$

$$\hat{w} = [-\sin \varphi \cos \vartheta, -\sin \varphi \sin \vartheta, \cos \varphi].$$

These unit vectors form a right-handed orthonormal triad attached to the spacecraft position, where: \hat{u} is radial, \hat{v} is tangential in the increasing longitude direction, and \hat{w} points toward increasing latitude.

The corresponding rotation matrix from the inertial frame to the ZEN frame is therefore:

$$\mathbf{R}(\vartheta, \varphi) = \begin{bmatrix} \hat{u}^T \\ \hat{v}^T \\ \hat{w}^T \end{bmatrix} = \begin{bmatrix} \cos \vartheta \cos \varphi & \sin \vartheta \cos \varphi & \sin \varphi \\ -\sin \vartheta & \cos \vartheta & 0 \\ -\cos \vartheta \sin \varphi & -\sin \vartheta \sin \varphi & \cos \varphi \end{bmatrix}. \quad (2.7)$$

Given the inertial components (x_b, y_b, z_b) of \mathbf{r}_{Eb} , the ZEN components are obtained through:

$$\mathbf{r}'_{Eb} = \mathbf{R}(\vartheta, \varphi) \mathbf{r}_{Eb},$$

which yields explicitly [9]:

$$(r_{Eb})_u = x_b \cos \vartheta \cos \varphi + y_b \sin \vartheta \cos \varphi + z_b \sin \varphi,$$

$$(r_{Eb})_v = -x_b \sin \vartheta + y_b \cos \vartheta,$$

$$(r_{Eb})_w = -x_b \cos \vartheta \sin \varphi - y_b \sin \vartheta \sin \varphi + z_b \cos \varphi.$$

In the ZEN formulation, the spacecraft position vector is written as:

$$\mathbf{r} = r \hat{u},$$

so that its coordinate representation in the local basis reads:

$$\mathbf{r}' = [r, 0, 0]^T.$$

The spacecraft-to-body relative vector in ZEN coordinates is therefore:

$$\mathbf{r}'_{bSC} = \mathbf{r}' - \mathbf{r}'_{Eb} = \begin{bmatrix} r - (r_{Eb})_u \\ -(r_{Eb})_v \\ -(r_{Eb})_w \end{bmatrix}.$$

The corresponding spacecraft-to-body distance is:

$$r_{bSC} = \sqrt{[r - (r_{Eb})_u]^2 + (r_{Eb})_v^2 + (r_{Eb})_w^2}.$$

Carrying out explicitly the projection of (2.6) onto the local ZEN basis, and exploiting the relations derived above for $(r_{Eb})_u$, $(r_{Eb})_v$ and $(r_{Eb})_w$, the third-body perturbing acceleration components can be written as:

$$\begin{aligned}(a_b)_u &= \mu_b \left[\frac{(r_{Eb})_u - r}{r_{bSC}^3} - \frac{(r_{Eb})_u}{r_{Eb}^3} \right], \\(a_b)_v &= \mu_b \left[\frac{(r_{Eb})_v}{r_{bSC}^3} - \frac{(r_{Eb})_v}{r_{Eb}^3} \right], \\(a_b)_w &= \mu_b \left[\frac{(r_{Eb})_w}{r_{bSC}^3} - \frac{(r_{Eb})_w}{r_{Eb}^3} \right].\end{aligned}$$

The total lunisolar perturbation entering the dynamical equations is obtained by summing the Sun and Moon contributions component-wise:

$$(a_{LS})_i = (a_s)_i + (a_\ell)_i, \quad i \in \{u, v, w\}.$$

2.3.2.3 Derivatives for the Costate Equations

For the indirect optimal control formulation, the costate dynamics require the partial derivatives of the perturbing accelerations with respect to the state variables. In the lunisolar case, the dependence on (r, ϑ, φ) arises both through the relative spacecraft-to-body distance and through the rotation defining the ZEN components of the third-body position vector.

Let \mathbf{r}_{Eb} denote the inertial position of the perturbing body b (Sun or Moon) with respect to the Earth center. Its components in the local ZEN basis are obtained through the rotation matrix $\mathbf{R}(\vartheta, \varphi)$ (2.7):

$$\mathbf{r}'_{Eb} = \mathbf{R}(\vartheta, \varphi) \mathbf{r}_{Eb} = \begin{bmatrix} (r_{Eb})_u \\ (r_{Eb})_v \\ (r_{Eb})_w \end{bmatrix}.$$

It is convenient to introduce the relative vector:

$$\mathbf{d} = \mathbf{r}'_{Eb} - \mathbf{r}' = \begin{bmatrix} (r_{Eb})_u - r \\ (r_{Eb})_v \\ (r_{Eb})_w \end{bmatrix},$$

together with the scalar quantities:

$$\rho_d = \|\mathbf{d}\|, \quad \rho_b = \|\mathbf{r}'_{Eb}\| = \|\mathbf{r}_{Eb}\|.$$

Here ρ_d represents the spacecraft-to-body distance, while ρ_b is the Earth-to-body distance (invariant under rotation).

With these definitions, the third-body acceleration components in the ZEN basis can be written compactly as:

$$(a_b)_i = \mu_b \left(\frac{d_i}{\rho_d^3} - \frac{(r_{Eb})_i}{\rho_b^3} \right), \quad i \in \{u, v, w\}.$$

Derivatives with respect to r . Since the rotated body components $(r_{Eb})_i$ do not depend on r , the only r -dependence arises through $d_u = (r_{Eb})_u - r$ and through ρ_d . Noting that:

$$\frac{\partial \mathbf{d}}{\partial r} = \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix},$$

and using the identity:

$$\frac{\partial}{\partial s} \left(\frac{v_i}{\rho^3} \right) = \frac{1}{\rho^3} \frac{\partial v_i}{\partial s} - 3 \frac{v_i}{\rho^5} \left(\mathbf{v} \cdot \frac{\partial \mathbf{v}}{\partial s} \right),$$

one obtains, for a single perturbing body b :

$$\begin{aligned} \frac{\partial (a_b)_u}{\partial r} &= \mu_b \left(-\frac{1}{\rho_d^3} + 3 \frac{d_u^2}{\rho_d^5} \right), \\ \frac{\partial (a_b)_v}{\partial r} &= 3\mu_b \frac{d_v d_u}{\rho_d^5}, \\ \frac{\partial (a_b)_w}{\partial r} &= 3\mu_b \frac{d_w d_u}{\rho_d^5}. \end{aligned}$$

The total radial derivatives entering the costate equations are obtained by summing the contributions of the Sun and the Moon.

Derivatives with respect to ϑ and φ . For $q \in \{\vartheta, \varphi\}$, the dependence arises entirely from the rotation matrix $\mathbf{R}(\vartheta, \varphi)$. In particular:

$$\frac{\partial \mathbf{r}'_{Eb}}{\partial q} = \frac{\partial \mathbf{R}}{\partial q} \mathbf{r}_{Eb}, \quad \frac{\partial \mathbf{d}}{\partial q} = \frac{\partial \mathbf{r}'_{Eb}}{\partial q}.$$

Applying again the vector identity above, the componentwise sensitivities can be written as:

$$\frac{\partial (a_b)_i}{\partial q} = \mu_b \left[\frac{1}{\rho_d^3} \frac{\partial d_i}{\partial q} - 3 \frac{d_i}{\rho_d^5} \left(\mathbf{d} \cdot \frac{\partial \mathbf{d}}{\partial q} \right) - \frac{1}{\rho_b^3} \frac{\partial (r_{Eb})_i}{\partial q} + 3 \frac{(r_{Eb})_i}{\rho_b^5} \left(\mathbf{r}'_{Eb} \cdot \frac{\partial \mathbf{r}'_{Eb}}{\partial q} \right) \right], \quad i \in \{u, v, w\}.$$

As in the geopotential case, the lunisolar perturbation depends only on the positional variables (r, ϑ, φ) ; therefore, no derivatives with respect to the velocity components (u, v, w) nor with respect to the spacecraft mass m arise from this contribution in the costate equations.

2.4 Complete Perturbed Equations of Motion

The complete dynamical model is obtained by augmenting the baseline spherical ZEN equations introduced in Section 2.2 with the perturbing acceleration terms derived in the previous sections.

The kinematic relations remain unchanged, as they follow directly from the geometric

definition of the ZEN velocity components. The dynamical equations are modified by including the total perturbing acceleration components $(a_p)_u$, $(a_p)_v$, and $(a_p)_w$, defined as:

$$(a_p)_i = (a_J)_i + (a_{LS})_i, \quad i \in \{u, v, w\}.$$

Accordingly, the perturbed equations of motion can be written as:

$$\begin{cases} \dot{r} = u, \\ \dot{\vartheta} = \frac{v}{r \cos \varphi}, \\ \dot{\varphi} = \frac{w}{r}, \\ \dot{u} = -\frac{\mu_E}{r^2} + \frac{v^2}{r} + \frac{w^2}{r} + \frac{T_u}{m} + (a_p)_u, \\ \dot{v} = -\frac{uv}{r} + \frac{vw}{r} \tan \varphi + \frac{T_v}{m} + (a_p)_v, \\ \dot{w} = -\frac{uw}{r} - \frac{v^2}{r} \tan \varphi + \frac{T_w}{m} + (a_p)_w, \\ \dot{m} = -\frac{T}{c}. \end{cases}$$

where μ_E denotes the Earth's gravitational parameter, $T = \|\mathbf{T}\|$ is the thrust magnitude, and c is the effective exhaust velocity. The quantities T_u , T_v , and T_w represent the projections of the thrust vector onto the local ZEN directions.

3 Conjunction Geometry and Probability of Collision

This chapter develops the collision-probability analysis that supports the definition of the displacement target to be enforced in the subsequent optimal-control problem. The objective is to characterize the encounter geometry at closest approach, to model the associated uncertainty in a consistent two-dimensional framework, and to evaluate the probability of collision in a form suitable for systematic target selection.

The analysis starts from a nominal propagation of the spacecraft and debris trajectories and proceeds with the identification of the Time of Closest Approach (TCA). Since the encounter geometry is highly sensitive to the timing of this event, a local refinement of the discrete minimum of the relative distance is introduced in order to improve the temporal accuracy of the TCA without modifying the global sampling step. The refined epoch defines the reference configuration at which the relative state and the uncertainty description are consistently evaluated.

At this configuration, a two-dimensional encounter plane orthogonal to the relative velocity is constructed, and the three-dimensional relative motion is projected onto it. The uncertainty at the encounter is represented in the same plane. In the absence of analytical covariance propagation from orbit determination, the planar mean and covariance are estimated empirically through Monte Carlo propagation of perturbed initial conditions up to the refined TCA, leading to a bivariate Gaussian description of the projected relative position. This statistical model provides the foundation for the subsequent collision-probability evaluation and admits a geometric interpretation through covariance ellipses.

Within this probabilistic framework, the probability of collision P_c is defined as the integral of the planar Gaussian density over a circular hard-body region. The same formulation is then used to evaluate P_c over a grid of candidate planar displacements at the TCA, generating probability maps that describe how different shifts modify the encounter risk. To ensure that the selected solution is not optimal only at a single instant, the probability is also evaluated on subsequent encounter planes, introducing a short-term robustness assessment. A discrete cost function is finally employed to identify a displacement that achieves the required probability reduction while limiting the magnitude of the shift. The resulting planar displacement is mapped back into three-dimensional space and adopted as the target condition for the optimal-control formulation addressed in the following chapters.

3.1 Short-Encounter Assumption

The collision-probability analysis developed in this chapter is formulated within the classical *short-encounter* framework, which provides a standard reduction of the three-dimensional conjunction problem to a two-dimensional representation in an encounter plane and leads to a computationally tractable definition of P_c [19, 20]. The short-encounter approximation is appropriate when the close approach occurs over a sufficiently short time interval around the Time of Closest Approach (TCA), so that the relative motion can be described locally with a simple kinematic model and the collision event can be assessed using the instantaneous geometry at the encounter epoch. In this setting, the assumptions listed below define the modeling hypotheses adopted throughout the chapter and clarify the scope of validity of the subsequent probability maps and target-selection procedure.

- **Locally rectilinear relative motion.** Over a sufficiently short time interval around the TCA, the relative trajectory can be approximated as linear with nearly constant relative velocity:

$$\boldsymbol{\rho}(t) \approx \boldsymbol{\rho}(t_{\text{TCA}}) + \mathbf{v}_{\text{rel}}(t_{\text{TCA}})(t - t_{\text{TCA}}),$$

so that the miss distance evolution is locally quadratic in time. This forms the basis of the 3D→2D reduction and of the parabolic TCA refinement.

- **Short encounter duration and frozen velocity effect.** Although uncertainty is present in both position and velocity at the initial epoch, the encounter duration is sufficiently short that the probabilistic analysis can be performed using the instantaneous distribution at the TCA. In this sense, velocity uncertainty contributes indirectly through the dispersion of the relative position at the TCA, while its dynamic effect during the encounter is neglected. The relative velocity is therefore treated as constant over the encounter window, consistently with the short-encounter formulation.
- **Position uncertainty frozen at TCA.** Because of the short encounter duration, the relative position uncertainty during the encounter is approximated by its value at the TCA. This allows the probabilistic analysis to be performed on the instantaneous projected distribution.
- **Collision governed by transverse geometry.** Within the locally linear model introduced above, the TCA corresponds to the instant at which the relative distance attains its minimum. In this approximation, the minimum-distance condition implies that the relative position at the TCA is orthogonal to the relative velocity:

$$\boldsymbol{\rho}_{\text{TCA}} \cdot \mathbf{v}_{\text{rel}}(t_{\text{TCA}}) = 0.$$

This property highlights that the minimum separation is governed by the components of the relative position orthogonal to \mathbf{v}_{rel} , whereas the component parallel to \mathbf{v}_{rel} primarily affects the timing of the encounter. As a consequence, the impact condition depends essentially on the transverse miss distance evaluated in the plane normal to \mathbf{v}_{rel} . The uncertainty component along \mathbf{v}_{rel} influences the encounter only through small timing shifts and therefore produces a second-order effect over the short encounter window. This justifies the reduction of the three-dimensional probabilistic problem to a two-dimensional integral performed in

the encounter plane orthogonal to \mathbf{v}_{rel} , consistently with standard short-term conjunction formulations [19, 20].

- **Gaussian uncertainty in the encounter plane.** The projected relative position is modeled as a bivariate Gaussian distribution characterized by mean and covariance $(\boldsymbol{\mu}, \mathbf{C})$, here obtained empirically from Monte Carlo propagation.
- **Hard-body representation.** Each object is modeled through an equivalent hard-body radius; collision occurs when the planar miss vector lies inside a disk of radius R , equal to the sum of the equivalent radii of the two objects.

These assumptions correspond to the standard short-term conjunction framework and are sufficient for the probabilistic model implemented in this work. They motivate, in particular, the emphasis on an accurate determination of the TCA, since the encounter plane and the projected geometry are evaluated at that epoch, the definition of an encounter plane orthogonal to $\mathbf{v}_{\text{rel}}(t_{\text{TCA}})$, and the reduction of the collision-probability computation to the integration of a bivariate Gaussian density over a hard-body disk in such plane.

3.2 Relative Motion and Time of Closest Approach

Consider two objects: a spacecraft (s) and a debris object (d). At time t their inertial position and velocity vectors are:

$$\mathbf{r}_s(t), \mathbf{v}_s(t) \in \mathbb{R}^3, \quad \mathbf{r}_d(t), \mathbf{v}_d(t) \in \mathbb{R}^3.$$

The relative quantities used throughout this chapter are defined as:

$$\boldsymbol{\rho}(t) := \mathbf{r}_d(t) - \mathbf{r}_s(t), \quad \mathbf{v}_{\text{rel}}(t) := \mathbf{v}_d(t) - \mathbf{v}_s(t).$$

The instantaneous separation distance is:

$$d(t) := \|\boldsymbol{\rho}(t)\|.$$

The *Time of Closest Approach* (TCA) corresponds to the time instant at which the relative distance $d(t)$ attains its minimum. From a geometric standpoint, the TCA identifies the configuration in which the two trajectories come closest in space. At this instant, the encounter geometry attains its minimum separation and the relative configuration used for the probabilistic analysis is defined. The encounter plane is constructed at the TCA, the relative motion is locally approximated in its vicinity, and the projected state at this epoch forms the basis of the subsequent collision-probability evaluation. Consequently, the accuracy of the TCA directly affects the encounter geometry and the computed value of the P_c .

3.2.1 Discrete Search

In the numerical implementation, the nominal propagation is available at discrete time instants and is therefore sampled on a temporal grid:

$$\{t_k\}_{k=0}^{K-1}, \quad t_k = t_0 + k \Delta t,$$

where k is an integer counter, t_0 is the initial time, and Δt is the sampling step.

Defining the sampled distances $d_k := \|\boldsymbol{\rho}(t_k)\|$, a first coarse estimate of the TCA is obtained by selecting the discrete minimum:

$$k^* := \arg \min_{k \in \{0, \dots, K-1\}} d_k, \quad t_{\text{TCA}} \approx t_{k^*}.$$

In analytical treatments of short-term encounters, such as the formulation proposed by Akella and Alfriend [19], the TCA is typically obtained by exploiting the linearized relative motion model, which allows a closed-form computation of the time of minimum separation under rectilinear and uniform motion assumptions. In the present work, however, the trajectories are generated through numerical propagation of the full dynamical model and are available only at sampled time instants. For this reason, the TCA is identified directly from the discrete propagation output rather than from an analytical minimization of a simplified relative-motion model.

The discrete estimate is affected by a quantization error of order $\mathcal{O}(\Delta t)$. While this may appear small in absolute terms, its impact on collision-probability assessment can be non-negligible. The encounter plane is constructed using the relative velocity at the TCA, and the relative position projected onto that plane defines the reference encounter configuration used in the probabilistic model. Even small timing offsets may therefore alter the orientation of the encounter plane, shift the projected miss vector, and modify the covariance obtained from Monte Carlo propagation. Since the probability of collision P_c depends sensitively on the relative placement of the Gaussian density with respect to the hard-body region, inaccuracies in the TCA may propagate into noticeable variations of the computed risk.

3.2.2 Parabolic Refinement

To improve the TCA estimate without globally reducing Δt , a local refinement is performed around the discrete minimum of the relative distance.

The discrete estimate t_{k^*} is sensitive to the sampling resolution. Reducing Δt globally would increase computational cost, especially when Monte Carlo propagation is involved. A more efficient strategy consists in locally reconstructing the behavior of the distance function in the neighborhood of the minimum, thereby achieving sub-step accuracy while preserving the original propagation grid.

The refinement operates on three consecutive samples around k^* :

$$(t_{k^*-1}, y_{k^*-1}), (t_{k^*}, y_{k^*}), (t_{k^*+1}, y_{k^*+1}), \quad y_k := d_k^2 = \|\boldsymbol{\rho}(t_k)\|^2.$$

The squared distance is used instead of d_k for both theoretical and numerical reasons. In the local linear regime around the TCA, the squared relative distance is a quadratic function of time, with its minimum located at the closest-approach epoch. Fitting a parabola to $d^2(t)$ is therefore fully consistent with the underlying encounter model and yields a sub-step estimate of the minimum without refining the global propagation grid.

From a numerical standpoint, $d^2(t)$ is smoother in the neighborhood of the minimum and avoids the additional nonlinearity introduced by the square root. Moreover, the curvature of $d^2(t)$ near the minimum is directly related to $\|v_{\text{rel}}\|^2$, which improves the conditioning of the interpolation problem and yields a stable estimate of the minimum location.

The three samples are interpolated with a parabola:

$$\widehat{y}(t) = at^2 + bt + c,$$

imposing $\widehat{y}(t_{k^*+j}) = y_{k^*+j}$ for $j \in \{-1, 0, +1\}$. The coefficients (a, b, c) are obtained by solving the corresponding 3×3 interpolation system.

The minimum of the interpolating parabola is obtained from the stationarity condition:

$$\frac{d\widehat{y}}{dt}(t) = 2at + b = 0, \quad t_{\text{TCA,ref}} = -\frac{b}{2a}.$$

If the interpolation becomes numerically degenerate (for instance due to a nearly singular system or $|a|$ too small), the algorithm falls back to the discrete estimate t_{k^*} . The refined time is finally constrained to the interval $[t_{k^*-1}, t_{k^*+1}]$ to avoid extrapolation and preserve consistency with the discrete sampling window.

Figure 3.1 illustrates the relative distance profile and the effect of the parabolic refinement. In the illustrated case, the discrete estimate yields $t_{k^*} = 1450.000$ s, while the refined value obtained from the quadratic interpolation is $t_{\text{TCA,ref}} = 1450.273019459666$ s, highlighting the sub-step correction introduced by the local refinement.

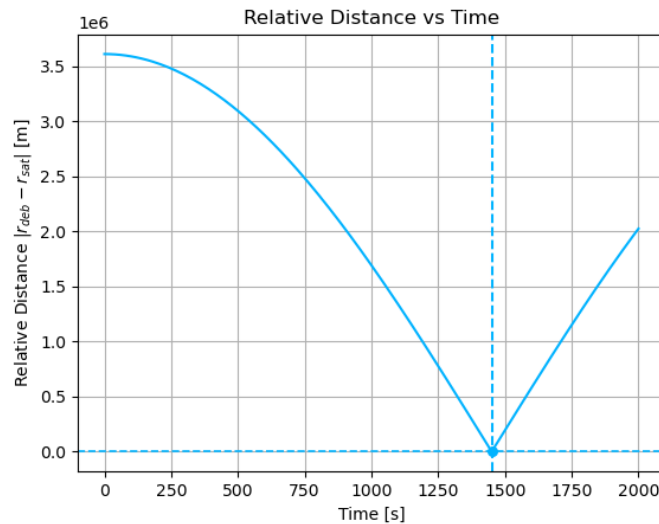


Figure 3.1: Relative distance profile and parabolic refinement of the time of closest approach (TCA).

3.3 Encounter Plane Definition

The evaluation of the two-dimensional collision probability relies on the definition of an *encounter plane*, onto which the three-dimensional relative motion is projected. The

reduction to such plane reflects the fact that, at closest approach, the collision event is governed by the transverse geometry with respect to the relative velocity. This allows the probabilistic problem to be formulated as a two-dimensional integral in a plane orthogonal to the relative velocity [19].

Let $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ be a right-handed orthonormal triad such that \mathbf{e}_3 is the normal to the encounter plane and $\mathbf{e}_1, \mathbf{e}_2$ span the plane itself. The planar projection of a vector $\mathbf{x} \in \mathbb{R}^3$ is defined as:

$$\Pi(\mathbf{x}) := \begin{bmatrix} \mathbf{e}_1^\top \mathbf{x} \\ \mathbf{e}_2^\top \mathbf{x} \end{bmatrix} \in \mathbb{R}^2.$$

The projected miss vector is therefore:

$$\mathbf{p}(t) := \Pi(\boldsymbol{\rho}(t)),$$

and its value at the refined TCA represents the mean of the probabilistic model used for collision assessment.

3.3.1 Plane Perpendicular to Relative Velocity

In standard conjunction assessment practice, the encounter plane is defined as perpendicular to the relative velocity at the TCA. At the instant of minimum separation, the relative position vector is orthogonal to the relative velocity in the locally linear model, so that the minimum distance is determined by the components transverse to \mathbf{v}_{rel} . The component parallel to \mathbf{v}_{rel} primarily affects the timing of the encounter, while the transverse components determine the impact condition. For this reason, the collision probability is evaluated in the plane orthogonal to \mathbf{v}_{rel} , where the relevant miss geometry is fully captured [19, 20].

Accordingly, the unit normal to the encounter plane is introduced as:

$$\mathbf{e}_3 := \frac{\mathbf{v}_{\text{rel}}(t_{\text{TCA,ref}})}{\|\mathbf{v}_{\text{rel}}(t_{\text{TCA,ref}})\|}.$$

This construction ensures orthogonality with respect to the instantaneous direction of relative motion at closest approach, thereby decoupling the along-track component from the two-dimensional probability integral.

A numerically robust construction of the remaining basis directions can be obtained by introducing the auxiliary vector:

$$\mathbf{a} := \mathbf{v}_d(t_{\text{TCA,ref}}) \times \mathbf{v}_s(t_{\text{TCA,ref}}),$$

which is orthogonal to both velocity vectors and encodes the instantaneous relative orbital geometry. The orthonormal triad is then completed as:

$$\mathbf{e}_1 := \frac{\mathbf{a} \times \mathbf{e}_3}{\|\mathbf{a} \times \mathbf{e}_3\|}, \quad \mathbf{e}_2 := \mathbf{e}_3 \times \mathbf{e}_1,$$

so that $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3)$ forms a right-handed orthonormal frame attached to the encounter configuration.

The specific orientation of $(\mathbf{e}_1, \mathbf{e}_2)$ within the plane does not influence the value of the collision probability, provided that both the covariance and the mean are expressed consistently in the same basis. What is physically relevant is the definition of the plane itself, namely its normal direction \mathbf{e}_3 . The above construction mirrors the geometric approach adopted in [19] while remaining fully consistent with the numerically propagated trajectories used in this work.

3.4 Gaussian Model in the Encounter Plane

3.4.1 Projection of State and Covariance

The construction of the Gaussian model in the encounter plane requires a consistent mapping from the initial state uncertainty to the planar relative geometry at the Time of Closest Approach (TCA). In the present work, this mapping is performed statistically through Monte Carlo propagation, coherently with the state parametrization adopted throughout the thesis.

Each object is described by a 7-dimensional state vector $x_0 \in \mathbb{R}^7$ (position, velocity, and mass) expressed in the spherical ZEN formulation. The use of the ZEN parametrization ensures consistency with the dynamical model introduced in the previous chapters and allows the uncertainty to be assigned directly to physically meaningful orbital components.

For the numerical examples reported in this thesis, the nominal initial states of the satellite and the debris are summarized in Table 3.1.

State Component	Satellite	Debris
r [m]	$R_E + 600.000$	$R_E + 600.050$
ϑ [rad]	0	0
φ [rad]	$\pi/6$	$\pi/3$
u [m/s]	0	0
v [m/s]	$\sqrt{\mu/r_s}$	$\sqrt{\mu/r_d}$
w [m/s]	0	0
m [kg]	462	1

Table 3.1: *Nominal initial states of satellite and debris.*

where R_E and μ denote the Earth's radius and gravitational parameter, respectively.

Uncertainty is introduced at the initial epoch through object-specific 1σ dispersions assigned independently to each state component. The adopted standard deviations are reported in Table 3.2.

1σ Dispersion	Satellite	Debris
σ_r [m]	30	60
σ_ϑ [rad]	5×10^{-6}	8×10^{-6}
σ_φ [rad]	5×10^{-6}	8×10^{-6}
σ_u [m/s]	5×10^{-3}	1×10^{-2}
σ_v [m/s]	5×10^{-3}	1×10^{-2}
σ_w [m/s]	5×10^{-3}	1×10^{-2}
σ_m [kg]	0	0

Table 3.2: 1σ uncertainties assigned to the initial state components.

In the absence of covariance matrices provided by an orbit determination process, the uncertainty propagation is performed numerically rather than analytically. This approach is consistent with the Monte Carlo procedures commonly adopted in conjunction analysis to characterize the relative position uncertainty at the closest point of approach [21]. Diagonal initial covariance matrices are therefore assumed for both objects:

$$\Sigma_{x_0}^s = \text{diag}\left((\sigma_s)^2\right), \quad \Sigma_{x_0}^d = \text{diag}\left((\sigma_d)^2\right).$$

The diagonal structure reflects the simplifying assumption of statistically independent state components at the initial epoch. Although real orbit determination solutions generally exhibit correlated uncertainties, the adopted model provides a controlled and transparent uncertainty scenario suitable for methodological analysis.

Independent Gaussian realizations are generated by perturbing each state component according to the corresponding standard deviation and propagating the ensemble to the refined TCA. This procedure naturally accounts for the nonlinear dynamics of the full propagation model and avoids relying on local linear covariance transport.

At the encounter epoch, each realization is projected into the encounter plane defined in Section 3.3.1, producing a set of planar samples $\mathbf{p}^{(i)} \in \mathbb{R}^2$. These projected samples constitute the empirical distribution of the relative position in the two-dimensional encounter geometry and represent the statistical basis of the Gaussian approximation introduced in the following subsection. The selected dispersions represent a plausible “debris less known than spacecraft” scenario within a simplified uncertainty framework. The debris is assigned approximately twice the radial dispersion (60 m vs. 30 m) and larger angular uncertainties, reflecting a lower orbit determination accuracy and a less precise knowledge of its dynamical state. Angular standard deviations of order 10^{-6} rad correspond, at the considered orbital radius, to cross-track and along-track position uncertainties of the order of several meters to a few tens of meters, consistent with the radial dispersion and with a short-term encounter analysis.

Velocity dispersions are set at the millimeter-per-second level (5–10 mm/s), sufficient to generate a noticeable spread at the TCA over typical propagation horizons while remaining within a conservative range for a preliminary assessment. This ensures that the resulting planar dispersion is driven both by positional and velocity uncertainty contributions accumulated during propagation.

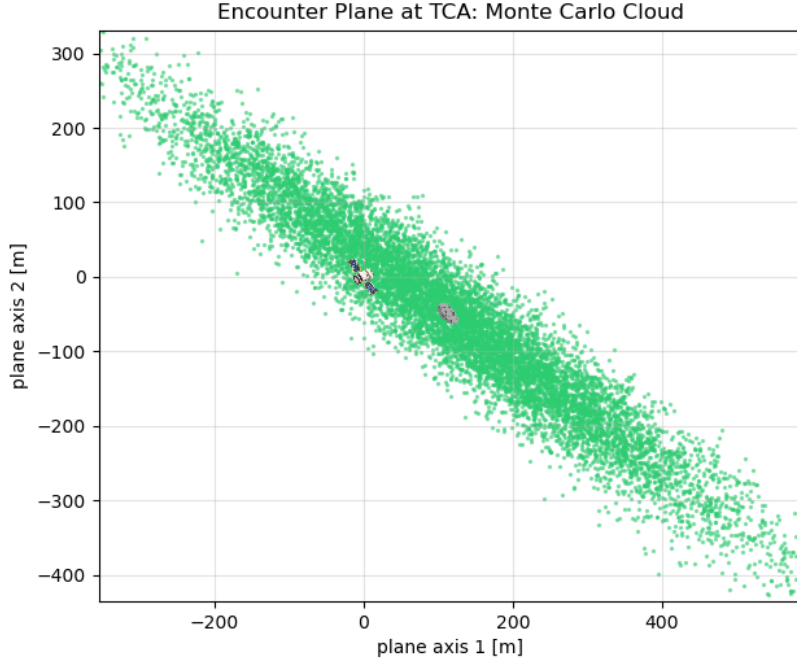


Figure 3.2: Monte Carlo cloud of the projected relative position at the TCA.

Finally, the mass component is treated as deterministic in the collision probability evaluation by setting $\sigma_m = 0$ for both objects. This assumption is consistent with the adopted two-dimensional short-encounter formulation, in which the collision probability depends on the relative position uncertainty at the encounter (and indirectly on velocity through the encounter plane definition), whereas the mass does not influence the kinematic projection and becomes relevant only in subsequent maneuver design and optimal control phases.

3.4.2 Mean and Covariance Estimation

Given the set of projected planar samples $\mathbf{p}^{(i)} \in \mathbb{R}^2$ obtained from the Monte Carlo propagation to the refined TCA, the statistical characterization of the encounter geometry is performed through the sample mean and covariance matrix [22]:

$$\boldsymbol{\mu} := \frac{1}{N} \sum_{i=1}^N \mathbf{p}^{(i)}, \quad \mathbf{C} := \frac{1}{N-1} \sum_{i=1}^N (\mathbf{p}^{(i)} - \boldsymbol{\mu})(\mathbf{p}^{(i)} - \boldsymbol{\mu})^\top. \quad (3.1)$$

The vector $\boldsymbol{\mu}$ represents the empirical mean relative position in the encounter plane at the TCA and can be interpreted as the most probable planar miss vector under the adopted uncertainty scenario. The covariance matrix $\mathbf{C} \in \mathbb{R}^{2 \times 2}$ quantifies the dispersion of the projected samples around $\boldsymbol{\mu}$, encoding not only the magnitude of the positional uncertainty but also its orientation and correlation structure within the encounter plane.

Since \mathbf{C} is symmetric positive definite, its eigendecomposition provides a direct geometric interpretation of the uncertainty region. The eigenvalues define the principal

variances along the major and minor axes of the uncertainty ellipse, while the corresponding eigenvectors determine its orientation with respect to the encounter-plane reference frame [23]. Together, $\boldsymbol{\mu}$ and \mathbf{C} provide a compact statistical description of the projected relative motion at the TCA.

Under the short-encounter assumption and for moderate uncertainty levels, the planar relative position is approximated as a bivariate Gaussian random vector [19, 22, 24]:

$$\mathbf{p} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C}), \quad f(\mathbf{p}) = \frac{1}{2\pi\sqrt{\det \mathbf{C}}} \exp\left(-\frac{1}{2}(\mathbf{p} - \boldsymbol{\mu})^\top \mathbf{C}^{-1}(\mathbf{p} - \boldsymbol{\mu})\right).$$

The use of a bivariate Gaussian model in the encounter plane is standard in short-encounter collision probability formulations and has been extensively validated through Monte Carlo analyses of relative position uncertainty at the closest approach [19, 21]. This Gaussian approximation provides a continuous probabilistic representation of the discrete Monte Carlo ensemble and enables a tractable formulation of the collision probability as the integral of the density over the hard-body region, while preserving the essential geometric features of the uncertainty distribution captured by the Monte Carlo samples.

The estimated pair $(\boldsymbol{\mu}, \mathbf{C})$ therefore constitutes the statistical description of the encounter configuration at the TCA and forms the probabilistic basis of the subsequent collision probability computation. The same representation is also employed to construct displacement maps and probability contours associated with the relative motion uncertainty.

Finally, the reliability of the estimated moments depends on the Monte Carlo sample size N . For sufficiently large N , the estimators in Eq. (3.1) converge to the corresponding mean and covariance of the projected distribution, ensuring consistency of the resulting collision probability evaluation [22].

3.4.3 Covariance Ellipse and Mahalanobis Distance

Covariance ellipses are the level sets of constant Mahalanobis distance associated with the bivariate Gaussian distribution describing the relative position in the encounter plane. For a Gaussian random vector with mean $\boldsymbol{\mu}$ and covariance \mathbf{C} , the probability density depends on the quadratic form $(\mathbf{p} - \boldsymbol{\mu})^\top \mathbf{C}^{-1}(\mathbf{p} - \boldsymbol{\mu})$, so that contours of constant density coincide with contours of constant Mahalanobis distance [23, 24]. These level sets are defined as:

$$(\mathbf{p} - \boldsymbol{\mu})^\top \mathbf{C}^{-1}(\mathbf{p} - \boldsymbol{\mu}) = k^2,$$

where k controls the Mahalanobis level and therefore the extent of the contour.

These ellipses provide a compact geometric description of the orientation, correlation, and dispersion of the positional uncertainty in the encounter plane.

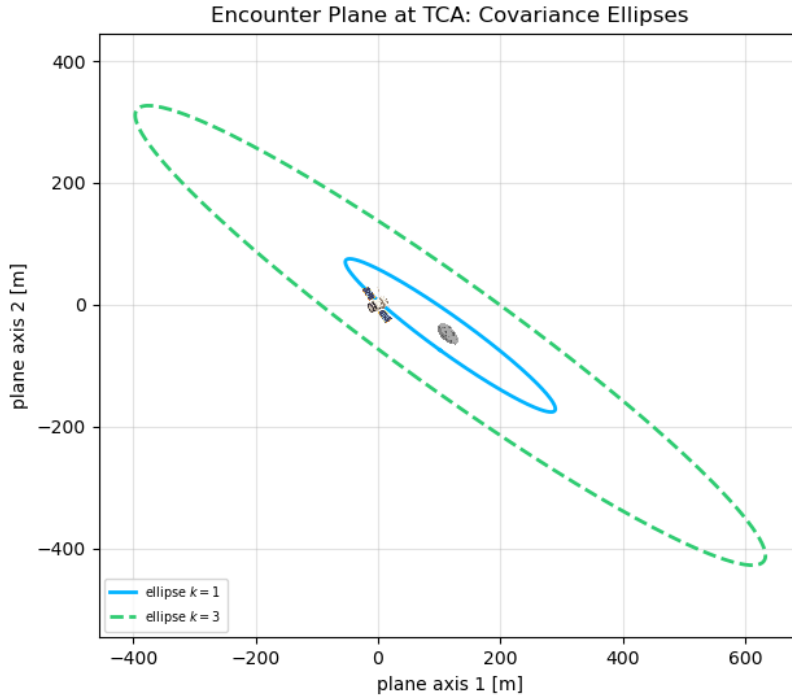


Figure 3.3: Covariance ellipses of the projected relative position distribution.

Since \mathbf{C} is symmetric positive definite, its eigenstructure fully determines the geometry of the ellipse: the principal directions coincide with the eigenvectors of \mathbf{C} , while the semi-axes are proportional to the square roots of its eigenvalues [23, 24].

Let the eigendecomposition of the covariance matrix be:

$$\mathbf{C} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^\top, \quad \mathbf{\Lambda} = \text{diag}(\lambda_1, \lambda_2), \quad \mathbf{V} = [\mathbf{v}_1 \ \mathbf{v}_2],$$

where $\mathbf{v}_1, \mathbf{v}_2$ form an orthonormal basis and $\lambda_1, \lambda_2 > 0$ are the principal variances. The inverse covariance is therefore:

$$\mathbf{C}^{-1} = \mathbf{V}\mathbf{\Lambda}^{-1}\mathbf{V}^\top, \quad \mathbf{\Lambda}^{-1} = \text{diag}\left(\frac{1}{\lambda_1}, \frac{1}{\lambda_2}\right).$$

Introducing principal coordinates aligned with the eigenvectors:

$$\boldsymbol{\xi} := \mathbf{V}^\top(\mathbf{p} - \boldsymbol{\mu}), \quad \boldsymbol{\xi} = \begin{bmatrix} \xi_1 \\ \xi_2 \end{bmatrix}, \quad (3.2)$$

the Mahalanobis condition reduces to:

$$\boldsymbol{\xi}^\top \mathbf{\Lambda}^{-1} \boldsymbol{\xi} = k^2,$$

which, written componentwise, yields:

$$\frac{\xi_1^2}{\lambda_1} + \frac{\xi_2^2}{\lambda_2} = k^2.$$

Dividing by k^2 , the canonical ellipse equation in principal coordinates is obtained:

$$\frac{\xi_1^2}{(k\sqrt{\lambda_1})^2} + \frac{\xi_2^2}{(k\sqrt{\lambda_2})^2} = 1, \quad (3.3)$$

showing that the semi-axes of the covariance ellipse are $k\sqrt{\lambda_1}$ and $k\sqrt{\lambda_2}$, aligned with \mathbf{v}_1 and \mathbf{v}_2 .

Geometrically, the eigendecomposition of \mathbf{C} induces an orthogonal change of basis that diagonalizes the quadratic form. In principal coordinates the Mahalanobis distance reduces to the sum of squared components weighted by the inverse principal variances, so that each coordinate is effectively normalized by the corresponding standard deviation [24].

A convenient parametric representation can be obtained by introducing an angular parameter $\theta \in [0, 2\pi)$ in principal coordinates and defining:

$$\xi_1(\theta) = k\sqrt{\lambda_1} \cos \theta, \quad \xi_2(\theta) = k\sqrt{\lambda_2} \sin \theta,$$

which identically satisfies (3.3). Transforming back to the original encounter-plane coordinates through (3.2), one obtains:

$$\mathbf{p}(\theta) = \boldsymbol{\mu} + k\sqrt{\lambda_1} \cos \theta \mathbf{v}_1 + k\sqrt{\lambda_2} \sin \theta \mathbf{v}_2, \quad \theta \in [0, 2\pi),$$

which describes an ellipse centered at $\boldsymbol{\mu}$, with principal axes aligned with the eigenvectors of the covariance matrix and semi-axis lengths equal to $k\sqrt{\lambda_1}$ and $k\sqrt{\lambda_2}$. This construction is fully consistent with the geometric interpretation of constant-density contours of a bivariate Gaussian distribution [23, 24].

From a probabilistic perspective, the squared Mahalanobis distance:

$$D^2 = (\mathbf{p} - \boldsymbol{\mu})^\top \mathbf{C}^{-1} (\mathbf{p} - \boldsymbol{\mu})$$

follows a chi-square distribution with two degrees of freedom for a bivariate Gaussian. Therefore, the parameter k can be directly related to a confidence level, with values obtained from standard χ^2 tables [23]. Typical values include $k = 1$ (approximately 39% probability), $k = \sqrt{2.30}$ (approximately 68%), and $k = \sqrt{5.99}$ (approximately 95%).

In the context of conjunction analysis, covariance ellipses provide a geometric representation of the positional uncertainty and its orientation in the encounter plane, complementing the probabilistic interpretation underlying the collision probability evaluation. They allow a direct visualization of how the uncertainty region intersects the hard-body disk and how candidate displacement targets modify the relative placement between the Gaussian density and the impact region.

In the present work, covariance ellipses corresponding to $k = 1$ and $k = 3$ are used in the graphical analysis of the encounter plane.

The choice $k = 1$ represents the unit Mahalanobis contour, enclosing approximately 39% of the probability mass for a bivariate Gaussian distribution. This ellipse highlights the

high-density core of the uncertainty region and provides an immediate visual indication of the most likely relative positions at TCA.

The value $k = 3$, on the other hand, corresponds to a probability content of approximately 98.9%, thus capturing almost the entire dispersion of the relative position uncertainty. The associated ellipse offers a conservative geometric envelope of the uncertainty region and is particularly useful for assessing the relative placement of candidate displacement points with respect to the bulk of the probability mass.

The combined use of $k = 1$ and $k = 3$ contours therefore enables a simultaneous representation of the central uncertainty core and of an extended confidence region, supporting the qualitative interpretation of the probability maps and the selection of displacement targets in the encounter plane.

3.4.4 Integration over the Hard-Body Disk

In the encounter plane, the collision event is modeled as the membership of the relative position vector \mathbf{p} in a circular hard-body region of radius R centered at the spacecraft:

$$\mathcal{D} := \{\mathbf{p} \in \mathbb{R}^2 : \|\mathbf{p}\| \leq R\}.$$

Within the short-encounter framework, the probability of collision is therefore obtained by integrating the bivariate Gaussian density over the hard-body disk:

$$P_c = \iint_{\mathcal{D}} f(\mathbf{p}) d\mathbf{p}, \quad (3.4)$$

which represents the standard geometric formulation adopted in two-dimensional conjunction analysis [19, 20]. In this formulation, the collision probability is interpreted as the probability that the stochastic relative position at the TCA lies inside the impact region defined by the hard-body disk.

From a physical standpoint, the disk \mathcal{D} represents the set of relative positions for which the distance between the two objects is smaller than the combined hard-body radius, implying geometric overlap and therefore collision. The reduction to a circular region follows from the use of equivalent spherical bodies, so that the collision condition depends only on the transverse separation in the encounter plane. The formulation assumes spherical equivalent bodies and neglects attitude, shape, and orientation effects, which is consistent with the level of abstraction of the short-term probabilistic model and with standard conjunction assessment practice [19].

In the present work, the equivalent hard-body radius is set to:

$$R = 20 \text{ m},$$

as a conservative assumption in the absence of precise information on the effective collision radii of the spacecraft and debris. This choice reflects a cautious risk assessment approach and ensures that the resulting probability of collision is not underestimated with respect to plausible object dimensions and potential safety margins.

The integral in Eq. (3.4) does not admit a closed-form solution in the general correlated case, since the Gaussian density cannot be separated in polar coordinates when the covariance matrix is not isotropic. It is therefore evaluated numerically in the subsequent sections. The formulation nonetheless provides the fundamental probabilistic definition from which all short-encounter P_c models derive and establishes the geometric basis of the collision-probability computation adopted in this work.

3.5 Probability Maps and Planar Displacement

3.5.1 Planar Displacement at the Time of Closest Approach

The effect of a deterministic spacecraft displacement at the time of closest approach (TCA) is analyzed in the encounter plane in order to quantify how planar shifts modify the collision probability. This formulation provides the basis for the construction of the probability maps presented later, where the dependence of P_c on candidate planar shifts is systematically explored.

Recalling the definition of the relative position (miss vector) $\boldsymbol{\rho}$ at the TCA, a deterministic three-dimensional spacecraft translation $\delta\mathbf{r}_s$ produces a shifted relative configuration:

$$\boldsymbol{\rho}'(t_{\text{TCA}}) = \mathbf{r}_d - (\mathbf{r}_s + \delta\mathbf{r}_s) = \boldsymbol{\rho} - \delta\mathbf{r}_s.$$

Projecting into the encounter plane through the projection operator $\Pi(\cdot)$ yields:

$$\mathbf{p}' = \Pi(\boldsymbol{\rho}') = \Pi(\boldsymbol{\rho}) - \Pi(\delta\mathbf{r}_s) = \mathbf{p} - \delta\mathbf{p}, \quad \delta\mathbf{p} := \Pi(\delta\mathbf{r}_s),$$

where \mathbf{p} and \mathbf{p}' represent the projected miss vector before and after the displacement, respectively.

A planar displacement $\delta\mathbf{p}$ therefore corresponds to a rigid translation of the hard-body collision region in the encounter plane, while leaving the probabilistic distribution of the relative position unchanged. This reflects the assumption that the displacement is deterministic and does not alter the uncertainty statistics evaluated at the TCA. In other words, the deterministic maneuver modifies the geometric alignment between the Gaussian uncertainty and the impact region, without altering the underlying covariance structure at the TCA. The collision event after the displacement is described by:

$$\|\mathbf{p}'\| \leq R \iff \|\mathbf{p} - \delta\mathbf{p}\| \leq R,$$

so that the probability of collision associated with $\delta\mathbf{p}$ becomes:

$$P_c(\delta\mathbf{p}) = \iint_{\|\mathbf{p} - \delta\mathbf{p}\| \leq R} f(\mathbf{p}) d\mathbf{p}. \quad (3.5)$$

Equation (3.5) shows that the effect of the planar displacement is exclusively a translation of the integration domain (hard-body disk) in the encounter plane, while the Gaussian density $f(\mathbf{p})$ remains unchanged. The dependence of P_c on $\delta\mathbf{p}$ is therefore purely geometric and reflects the relative placement of the uncertainty distribution with respect to the collision region.

For numerical evaluation, it is convenient to introduce polar coordinates centered at δ_p , since the integration domain is circular. A generic point in the hard-body disk can be parameterized as:

$$\mathbf{p} = \delta_p + \begin{bmatrix} r \cos \theta \\ r \sin \theta \end{bmatrix}, \quad 0 \leq r \leq R, \quad 0 \leq \theta < 2\pi,$$

for which the condition $\|\mathbf{p} - \delta_p\| \leq R$ is automatically satisfied and the area element becomes:

$$d\mathbf{p} = r dr d\theta.$$

Substituting into Eq. (3.5) yields the polar form:

$$P_c(\delta_p) = \int_0^{2\pi} \int_0^R f\left(\delta_p + \begin{bmatrix} r \cos \theta \\ r \sin \theta \end{bmatrix}\right) r dr d\theta, \quad (3.6)$$

which is the expression used for the numerical evaluation of the collision probability and forms the computational core of the probability maps presented in the following section.

3.5.2 Numerical Quadrature and Construction of the Probability Map

In the adopted model, the projected relative position is assumed to follow a bivariate Gaussian distribution:

$$\mathbf{p} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C}),$$

where the mean $\boldsymbol{\mu}$ and covariance \mathbf{C} are estimated via Monte Carlo propagation of the initial state uncertainty and subsequent projection into the encounter plane, as discussed in Section 3.4. Under this assumption, the collision probability associated with a planar displacement δ_p is fully determined by the geometric alignment between the Gaussian density and the shifted hard-body disk.

The collision probability corresponding to a given planar displacement δ_p is evaluated by integrating the Gaussian density over the disk of radius R centered at δ_p , according to Eq. (3.5), or equivalently Eq. (3.6). Since no closed-form solution exists in the general correlated case, the integral is computed numerically by quadrature in polar coordinates.

From an operational standpoint, the construction of the probability map proceeds as follows:

1. A regular two-dimensional grid of planar displacements is defined:

$$\delta_p = (\delta_{p,i}, \delta_{p,j}),$$

covering the region of interest in the encounter plane.

2. For each grid node, the collision probability $P_c(\delta_p)$ is computed by numerical quadrature of Eq. (3.6) over the polar variables (r, θ) . This corresponds to evaluating the integral of the Gaussian density over the shifted hard-body disk for that specific displacement.

3. The resulting values are assembled into a two-dimensional scalar field:

$$P_c : \mathbb{R}^2 \rightarrow \mathbb{R}, \quad \delta_p \mapsto P_c(\delta_p),$$

which associates each planar displacement with the corresponding collision probability.

The resulting probability map provides a geometric and probabilistic representation of how different planar displacements modify the collision risk in the encounter plane. Regions of high probability correspond to configurations where the Gaussian uncertainty significantly overlaps the hard-body disk, whereas lower values indicate improved geometric separation between the uncertainty core and the collision region.

Contour levels are included to highlight iso-probability curves, facilitating the identification of displacement regions that achieve prescribed risk-reduction thresholds. These maps form the basis for the subsequent selection of candidate planar targets to be enforced by the optimal-control maneuver.

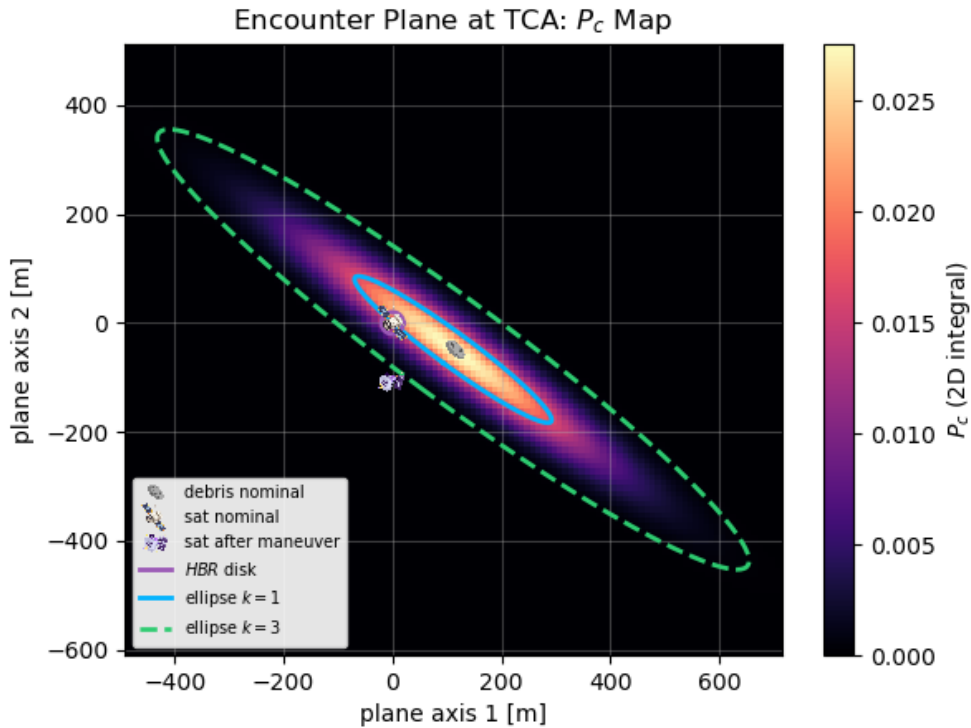


Figure 3.4: Collision probability map $P_c(\delta_p)$ over the displacement grid.

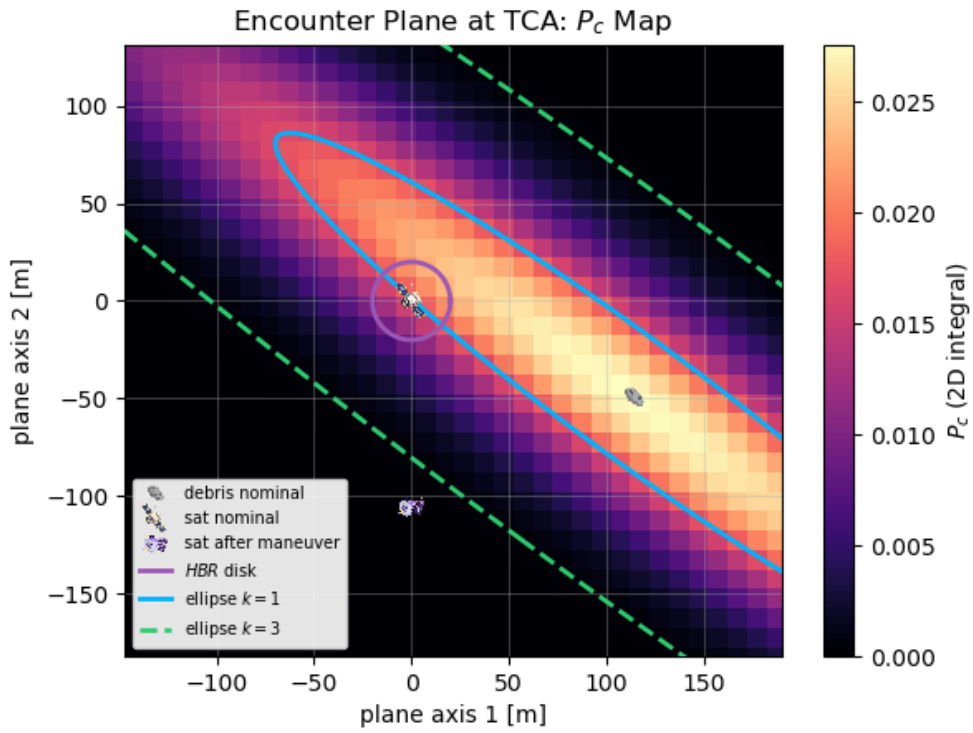


Figure 3.5: Zoomed view of the probability map in the region of interest.




Figure 3.6: Collision probability map with contour levels highlighting regions of decreasing P_c as the planar displacement varies.

3.5.3 Subsequent Encounter Planes and Robustness

While the primary constraint is enforced on the encounter plane at the refined TCA, the implementation also evaluates the collision probability on subsequent encounter planes in order to promote robustness over a finite time window following closest approach. This extension mitigates the sensitivity of the solution to the exact timing of the TCA and provides a short-term temporal consistency check of the selected displacement.

Specifically, in addition to the TCA plane ($t = t_{\text{TCA}}$), two further planes are considered at:

$$t_1 = t_{\text{TCA}} + 60 \text{ s}, \quad t_2 = t_{\text{TCA}} + 120 \text{ s}.$$

These offsets provide a practical robustness check over a short-term horizon in which the short-encounter assumptions remain approximately valid, while being sufficiently long to reveal whether a displacement that is optimal at the TCA rapidly loses effectiveness as the relative geometry evolves.

Denoting by $P_c^{(0)}(\delta_p)$ the probability evaluated on the TCA plane and by $P_c^{(1)}(\delta_p)$, $P_c^{(2)}(\delta_p)$ the probabilities on the planes at t_1 and t_2 , respectively, the resulting set of maps offers a compact description of the temporal sensitivity of a given planar displacement beyond the instantaneous closest approach. At each considered epoch, the relative position and velocity are recomputed from the propagated trajectories, the corresponding encounter plane is constructed, and the projected uncertainty distribution is evaluated consistently in that plane.

Within this geometric framework, the planar displacement δ_p is defined at the TCA as the target shift in the encounter plane, without yet specifying the maneuver required to realize it. The determination of the actual maneuver that produces such displacement is addressed later through the optimal-control formulation, where the selected planar shift is imposed as a terminal condition. Consequently, when evaluating $P_c^{(1)}(\delta_p)$ and $P_c^{(2)}(\delta_p)$, the same planar displacement identified at the TCA is consistently projected onto the subsequent encounter planes, without re-propagating a modified post-maneuver trajectory. The resulting assessment therefore provides a geometric robustness indicator rather than a fully dynamical one. It reflects how the natural evolution of the relative motion interacts with a displacement that achieves the desired separation at the TCA, highlighting configurations whose effectiveness is intrinsically sensitive to the short-term drift of the projected miss vector.

The relative motion induces a systematic evolution of the projected miss vector across successive planes, which can be interpreted geometrically as a *drift* direction in the encounter plane. After the TCA, the relative trajectory continues to evolve, and the projection of the relative position generally translates across the plane according to the local relative velocity. As a consequence, the geometric separation achieved by a displacement at the TCA may be partially reduced or reinforced depending on its alignment with this natural drift direction.

This behavior can be understood with a simple geometric argument. If the debris projected trajectory crosses the encounter plane from right to left, the projected miss vector will continue drifting leftward after the TCA. A displacement applied at the

TCA in the same direction (right to left) may reduce the probability at t_{TCA} by shifting the hard-body region away from the peak of the Gaussian density; however, as time advances and the miss vector continues to drift leftward, the displaced configuration may progressively lose separation and can lead to comparatively higher probabilities on subsequent planes. Conversely, displacements orthogonal to the drift or opposed to it tend to preserve geometric separation over a longer interval, yielding a more robust reduction of P_c across the considered time window.

This robustness effect is captured explicitly in the target selection cost introduced in Section 3.5.4, where candidate displacements are not ranked solely on $P_c^{(0)}$ but also on their performance on $P_c^{(1)}$ and $P_c^{(2)}$. In this way, the selected displacement is required to achieve not only instantaneous risk reduction but also short-term temporal stability of the encounter geometry.

The encounter geometry on the subsequent planes considered for robustness assessment is illustrated in Fig. 3.7, where the projected miss vector and the associated covariance ellipses are shown at two time instants after the TCA.

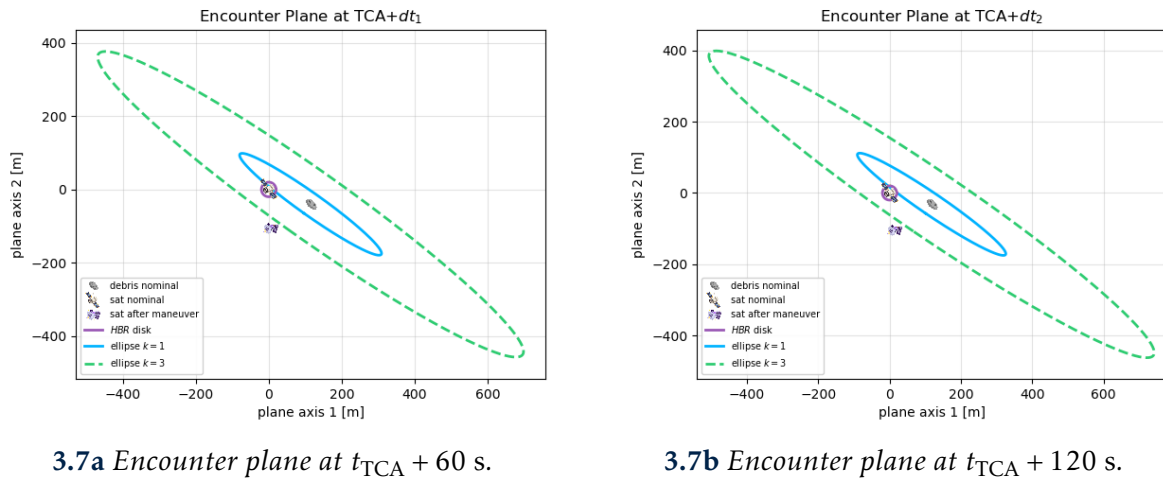


Figure 3.7: Projected encounter geometry on the subsequent planes used to assess temporal robustness.

3.5.4 Target Displacement Selection

Once the probability map $P_c(\delta_p)$ has been evaluated over the displacement grid, the final target displacement is selected by solving a discrete optimization problem defined over the same grid. The objective is to identify a planar shift that reduces the collision probability below an acceptable threshold while keeping the maneuver magnitude limited and ensuring robustness over the subsequent encounter planes.

A preliminary feasibility filter is first applied at the TCA plane:

$$P_c^{(0)}(\delta_p) \leq P_c^{\text{th}},$$

where P_c^{th} is a prescribed collision probability threshold, set here to 10^{-4} .

This constraint excludes candidate displacements that do not guarantee a sufficient level

of instantaneous risk reduction at closest approach. The adopted threshold is consistent with common conjunction assessment practice and represents a conservative safety requirement at the TCA, ensuring that only displacements achieving a meaningful probability decrease are considered for further ranking.

Among the feasible candidates, the optimal displacement is selected through the minimization of a scalar cost function:

$$J(\delta_p) = w_{\text{shift}} \|\delta_p\| + w_0 P_c^{(0)}(\delta_p) + w_1 P_c^{(1)}(\delta_p) + w_2 P_c^{(2)}(\delta_p) + w_{\text{align}} \max(0, \cos \alpha),$$

where $P_c^{(0)}$, $P_c^{(1)}$ and $P_c^{(2)}$ denote the collision probabilities evaluated at the TCA plane and at the two subsequent planes, respectively.

The term $\|\delta_p\|$ represents the magnitude of the planar displacement and provides a first-order proxy for maneuver effort, since larger geometric separations generally require larger velocity increments in the subsequent optimal-control problem. The probability terms $P_c^{(0)}$, $P_c^{(1)}$, and $P_c^{(2)}$ account for risk reduction at the TCA and for short-term temporal robustness, preventing the selection of displacements that are optimal only instantaneously but rapidly lose effectiveness as the encounter geometry evolves.

The alignment term depends on the angle α between the displacement δ_p and the nominal drift direction \mathbf{d} of the projected miss vector:

$$\cos \alpha = \frac{\delta_p \cdot \mathbf{d}}{\|\delta_p\| \|\mathbf{d}\|}.$$

The drift direction \mathbf{d} is defined geometrically as the finite difference of the projected relative position between the TCA plane and the second subsequent encounter plane. Denoting by $\mathbf{p}^{(0)}$ the projected miss vector at the TCA and by $\mathbf{p}^{(2)}$ the projected miss vector at $t_2 = t_{\text{TCA}} + 120$ s, the drift vector is computed as:

$$\mathbf{d} = \mathbf{p}^{(2)} - \mathbf{p}^{(0)}.$$

This construction provides a first-order approximation of the short-term geometric evolution of the projected miss vector in the encounter plane. Using the second subsequent plane yields a more representative estimate of the natural drift direction, reducing sensitivity to local fluctuations at the first offset and remaining consistent with the locally linear relative motion model adopted in the short-encounter framework.

The contribution $w_{\text{align}} \max(0, \cos \alpha)$ penalizes displacements aligned with this drift direction, since such maneuvers tend to lose geometric separation as the projected miss vector continues to evolve in the same direction. Conversely, transverse or opposite displacements are not penalized, as they are generally more effective in preserving separation over the considered time window and thus provide improved robustness.

In the present implementation, the weights are selected as:

$$w_{\text{shift}} = 1.0, \quad w_0 = 1.0, \quad w_1 = 1.0, \quad w_2 = 1.0, \quad w_{\text{align}} = 0.2.$$

This choice assigns comparable importance to maneuver magnitude and collision probability reduction across the three encounter planes, while introducing a moderate penalty for drift-aligned displacements. The relatively smaller value of w_{align} ensures that alignment acts as a secondary robustness criterion, influencing the selection primarily when candidate solutions exhibit similar probability reduction and displacement magnitude.

The optimal planar displacement is finally selected as:

$$\delta_p^* = \arg \min_{\delta_p} J(\delta_p),$$

within the feasible set defined by the threshold constraint. The selected displacement is subsequently mapped back into three-dimensional space, consistently with the adopted encounter plane basis, and is imposed as the terminal geometric target in the optimal-control problem developed in the following chapters.

In the present case, the selected displacement corresponds in the encounter plane to:

$$(\text{axis}_1, \text{axis}_2) = (-0.001, -105.000) \text{ m},$$

which maps, in the spherical ZEN frame, to a radial displacement of 105 m. This maneuver reduces the collision probability at the TCA from:

$$P_c = 1.701004 \times 10^{-2}$$

to:

$$P_c^{(0)} = 6.767695 \times 10^{-5},$$

thus satisfying the imposed safety threshold. The corresponding probabilities on the subsequent encounter planes are:

$$P_c^{(1)} = 4.392757 \times 10^{-5}, \quad P_c^{(2)} = 2.668747 \times 10^{-5}.$$

The probability not only meets the required constraint at the TCA but continues to decrease over the subsequent planes, indicating that the selected displacement remains well aligned with the short-term evolution of the relative geometry. This monotonic reduction confirms the robustness of the chosen target and validates the effectiveness of the combined probability and alignment criteria embedded in the cost function.

3.6 Monte Carlo Estimate of $P_c(\delta_p)$ (Validation)

As an independent verification of the numerical quadrature, the collision probability $P_c(\delta_p)$ can be estimated using a hit-or-miss Monte Carlo approach. Given a set of planar samples $\mathbf{p}^{(i)}$ drawn from the bivariate Gaussian distribution:

$$\mathbf{p}^{(i)} \sim \mathcal{N}(\boldsymbol{\mu}, \mathbf{C}),$$

the empirical estimate of the collision probability associated with a planar displacement δ_p is defined as:

$$\widehat{P}_c(\delta_p) = \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{\|\mathbf{p}^{(i)} - \delta_p\| \leq R\}, \quad (3.7)$$

where $\mathbf{1}\{\cdot\}$ denotes the indicator function of the collision event, equal to one when the shifted sample lies inside the hard-body disk and zero otherwise.

This estimator directly reflects the geometric definition of the collision event in the encounter plane and provides a statistical validation of the quadrature-based computation of $P_c(\delta_p)$. The estimator in Eq. (3.7) is unbiased and converges to the true probability as $N \rightarrow \infty$. For sufficiently large sample sizes, the Monte Carlo estimate therefore provides an independent numerical approximation of $P_c(\delta_p)$, enabling a direct consistency check of the quadrature-based integration procedure.

The use of a hit-or-miss formulation is consistent with standard validation practices in collision probability assessment and with the geometric interpretation underlying short-encounter models [20].

In the present implementation, the number of Monte Carlo samples N is limited by the available computational memory, which constrains the maximum number of realizations that can be stored and processed simultaneously. Despite this practical limitation, the comparison between the quadrature-based probability and the Monte Carlo estimate shows very small discrepancies. The mean absolute error over the displacement grid is:

$$\bar{\varepsilon} = 2.287 \times 10^{-7},$$

while the maximum observed error is:

$$\varepsilon_{\max} = 1.707 \times 10^{-5}.$$

These values are several orders of magnitude smaller than the collision probability threshold adopted for target selection (10^{-4}) and therefore confirm the numerical consistency of the quadrature scheme.

Although even smaller errors could in principle be achieved by increasing the number of Monte Carlo samples, the obtained accuracy is already well within the level required for conjunction assessment and robust target definition. The validation thus supports the reliability of the quadrature-based computation of $P_c(\delta_p)$ used throughout the displacement analysis.

4 Optimal Control Problem Formulation

The probabilistic analysis developed in the previous chapter has identified a geometrically admissible target displacement in the encounter plane, selected to guarantee a prescribed reduction of the collision probability. That displacement defines the desired relative configuration at the time of closest approach, but it does not specify how the spacecraft should maneuver to achieve it. A dynamically consistent trajectory connecting the initial state to the required terminal configuration must therefore be determined.

The spacecraft guidance problem is therefore cast as an OCP, in which the trajectory dynamics, boundary conditions, and performance objective are treated in a unified variational framework. The objective is to determine a thrust history that satisfies the equations of motion, enforces the prescribed terminal constraints, and minimizes propellant consumption under the adopted dynamical model.

Trajectory optimization problems of this kind are classically formulated within the Bolza framework and have been extensively studied in the context of spaceflight mechanics. From the primer-vector theory introduced by Lawden [25] to systematic treatments of indirect methods in standard references such as Kirk [26], Bryson and Ho [2], and Burghes and Graham [27], optimal control theory provides the mathematical structure required to characterize admissible optimal solutions. In this thesis, the formulation, notation, and modelling conventions are directly derived from the framework presented in [9, 10], which is adopted consistently throughout the chapter.

Particular attention is devoted to the structure of the fuel-optimal problem. Since thrust magnitude appears linearly in the Hamiltonian associated with the variational formulation, the optimal control is expected to consist of alternating thrust and coast arcs. This behaviour is governed by a switching function that compares the marginal benefit of thrusting with the associated cost in propellant mass, and it plays a central role in the indirect solution strategy adopted in this work.

Within this framework, the spacecraft dynamics are expressed in spherical ZEN coordinates, consistently with the adopted dynamical model. The objective function is defined in a Mayer form corresponding to final-mass maximization, while the boundary conditions are formulated in compact constraint notation. The resulting optimality conditions define a coupled state-costate system with endpoint constraints, providing the analytical structure required to determine a dynamically feasible and fuel-optimal collision-avoidance maneuver that achieves the selected geometric displacement.

4.1 General structure of the optimal control problem

An OCP is concerned with determining a control law that extremizes a prescribed performance measure while ensuring that the system dynamics and all imposed constraints are satisfied. In this context, the objective is not merely to find a trajectory connecting two configurations, but to identify, among all admissible trajectories, the one that optimally fulfills a given criterion.

The system evolution is described through a set of first-order ordinary differential equations governing the time variation of the state variables. Let the state vector be denoted by:

$$\mathbf{x}(t) \in \mathbb{R}^n,$$

where n represents the number of state components. The state evolution depends on a control vector:

$$\mathbf{u}(t) \in \mathbb{R}^m,$$

which contains m control variables and acts over a finite time interval $t \in [t_0, t_f]$, with t being the independent variable.

The dynamical model can therefore be written in compact form as:

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

where the function $\mathbf{f}(\cdot)$ defines how the state evolves as a function of the current state, the applied control, and possibly time.

An optimal solution consists of determining a trajectory $\mathbf{x}^*(t)$ and an associated control history $\mathbf{u}^*(t)$ such that the dynamical equations (4.1) are satisfied and the selected merit index attains an extremal value.

In general, the trajectory is constrained by boundary conditions imposed at the initial and final instants. The times t_0 and t_f represent the external boundaries of the problem. Constraints may involve state components at the initial and/or final time, and in some cases the times themselves may be fixed or subject to optimization.

All boundary conditions are conveniently collected in the compact form [9]:

$$\mathcal{X}(\mathbf{x}_0, \mathbf{x}_f, t_0, t_f) = \mathbf{0}, \quad (4.2)$$

where $\mathbf{x}_0 = \mathbf{x}(t_0)$ and $\mathbf{x}_f = \mathbf{x}(t_f)$, and the constraint function [9]:

$$\mathcal{X} : \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^q$$

contains the complete set of q independent boundary constraints.

In addition to endpoint constraints, the control vector itself may be subject to admissibility requirements. Certain components of $\mathbf{u}(t)$ can be restricted to lie within prescribed bounds, for instance through minimum and maximum values. The set of admissible controls is denoted as [9]:

$$\mathbf{u} \in \mathcal{U},$$

and defines the domain within which the optimization must be performed.

The problem thus takes the form of a two-point boundary value problem, where the system of differential equations is integrated between two extremal states under the boundary conditions (4.2), and the control must belong to the admissible set \mathcal{U} . A trajectory satisfying both the dynamical model and all constraints is considered admissible. Among these admissible trajectories, the optimal one is selected according to the chosen merit index, whose formulation is introduced in the following section.

4.1.1 Merit index and augmented formulation

The optimality of the control problem is assessed through the definition of a scalar merit index whose extremal value must be determined. Such a functional quantifies the performance of a given trajectory and allows one to discriminate, among all admissible solutions, the one that satisfies the prescribed optimality criterion.

In its most general representation, the merit index is written in Bolza form as [9]:

$$\mathcal{J} = \varphi(\mathbf{x}_0, \mathbf{x}_f, t_0, t_f) + \int_{t_0}^{t_f} \Phi(\mathbf{x}(t), \mathbf{u}(t), t) dt,$$

where the endpoint term $\varphi(\cdot)$ depends on the state and possibly time at the extremal boundaries, while the integral term $\Phi(\cdot)$ accounts for the cumulative contribution of state and control variables during the trajectory evolution.

The two contributions play conceptually distinct roles. The terminal function φ characterizes the final configuration reached by the system within the admissible set and therefore reflects the objective associated with the extremal states. Conversely, the running term Φ measures how the trajectory evolves between t_0 and t_f , capturing the influence of intermediate control actions and state variations over time.

For notational compactness, the following definitions are introduced [9]:

$$\varphi \triangleq \varphi(\mathbf{x}_0, \mathbf{x}_f, t_0, t_f), \quad \mathcal{X} \triangleq \mathcal{X}(\mathbf{x}_0, \mathbf{x}_f, t_0, t_f).$$

In order to derive the necessary conditions of optimality, the boundary constraints and the dynamical equations must be incorporated into the optimization process. This is achieved by introducing a set of adjoint variables $\lambda(t)$, associated with the state equations, and a vector of constant multipliers $\boldsymbol{\mu}$, associated with the boundary conditions.

The augmented merit index is therefore defined as [9]:

$$\mathcal{J}^* = \varphi + \boldsymbol{\mu}^\top \mathcal{X} + \int_{t_0}^{t_f} [\Phi + \boldsymbol{\lambda}^\top (\mathbf{f} - \dot{\mathbf{x}})] dt. \quad (4.3)$$

This construction ensures that any admissible solution of the original problem corresponds to a stationary point of the augmented functional. In particular, when the boundary conditions are satisfied ($\mathcal{X} = \mathbf{0}$) and the state equations hold ($\dot{\mathbf{x}} = \mathbf{f}$), the augmented functional reduces to the original merit index, that is:

$$\mathcal{J}^* = \mathcal{J}.$$

Hence, solving the augmented problem is mathematically equivalent to solving the original one, provided that all constraints are fulfilled.

The term involving $\lambda^\top \dot{\mathbf{x}}$ can be eliminated by integrating by parts. Since the state time derivatives are not known a priori, it is convenient to rewrite [9]:

$$\int_{t_0}^{t_f} -\lambda^\top \dot{\mathbf{x}} dt = -\lambda^\top \mathbf{x} \Big|_{t_0}^{t_f} + \int_{t_0}^{t_f} \dot{\lambda}^\top \mathbf{x} dt.$$

Substituting into (4.3) yields [9]:

$$\mathcal{J}^* = \varphi + \boldsymbol{\mu}^\top \boldsymbol{\mathcal{X}} + \lambda^\top(t_0) \mathbf{x}_0 - \lambda^\top(t_f) \mathbf{x}_f + \int_{t_0}^{t_f} (\boldsymbol{\Phi} + \lambda^\top \mathbf{f} - \dot{\lambda}^\top \mathbf{x}) dt.$$

A central quantity naturally arises in this formulation, namely the Hamiltonian function [9]:

$$\mathcal{H}(\mathbf{x}, \mathbf{u}, \boldsymbol{\lambda}, t) = \boldsymbol{\Phi}(\mathbf{x}, \mathbf{u}, t) + \lambda^\top \mathbf{f}(\mathbf{x}, \mathbf{u}, t).$$

The necessary condition for optimality requires that the augmented functional be stationary at the optimal solution. Therefore, its first variation must vanish. By computing the first-order variation of \mathcal{J}^* with respect to the independent variations δt_0 , δt_f , $\delta \mathbf{x}_0$, $\delta \mathbf{x}_f$, $\delta \mathbf{x}(t)$ and $\delta \mathbf{u}(t)$, one obtains [9]:

$$\begin{aligned} \delta \mathcal{J}^* = & \left(\frac{\partial \varphi}{\partial t_0} + \boldsymbol{\mu}^\top \frac{\partial \boldsymbol{\mathcal{X}}}{\partial t_0} - \mathcal{H}_0 \right) \delta t_0 + \\ & \left(\frac{\partial \varphi}{\partial t_f} + \boldsymbol{\mu}^\top \frac{\partial \boldsymbol{\mathcal{X}}}{\partial t_f} + \mathcal{H}_f \right) \delta t_f + \\ & \left(\frac{\partial \varphi}{\partial \mathbf{x}_0} + \boldsymbol{\mu}^\top \frac{\partial \boldsymbol{\mathcal{X}}}{\partial \mathbf{x}_0} + \lambda_0^\top \right) \delta \mathbf{x}_0 + \\ & \left(\frac{\partial \varphi}{\partial \mathbf{x}_f} + \boldsymbol{\mu}^\top \frac{\partial \boldsymbol{\mathcal{X}}}{\partial \mathbf{x}_f} - \lambda_f^\top \right) \delta \mathbf{x}_f + \\ & \int_{t_0}^{t_f} \left[\left(\frac{\partial \mathcal{H}}{\partial \mathbf{x}} + \dot{\boldsymbol{\lambda}} \right)^\top \delta \mathbf{x} + \frac{\partial \mathcal{H}}{\partial \mathbf{u}} \delta \mathbf{u} \right] dt. \end{aligned}$$

The arbitrariness of the independent variations implies that each multiplying coefficient must vanish. From the coefficients associated with δt_0 , δt_f , $\delta \mathbf{x}_0$, $\delta \mathbf{x}_f$, one obtains algebraic conditions at the extremal points, known respectively as transversality conditions (for time) and optimality conditions at the boundaries (for state variables).

From the integral term, two distinct sets of relations emerge. By nullifying the coefficient of $\delta \mathbf{x}$, one derives the differential equations governing the evolution of the adjoint variables. By nullifying the coefficient of $\delta \mathbf{u}$, one obtains a set of m algebraic equations for the control.

4.1.2 Boundary conditions for optimality

The optimality conditions associated with the extremal points of the trajectory arise from the coefficients of the independent variations of time and state at t_0 and t_f in the

first variation of the augmented functional. By enforcing the stationarity of \mathcal{J}^* with respect to δt_0 , δt_f , $\delta \mathbf{x}_0$, and $\delta \mathbf{x}_f$, a set of algebraic relations is obtained that prescribes how the solution must behave at the trajectory boundaries.

These relations naturally separate into two groups. The first group concerns variations of the initial and final times and is commonly referred to as the transversality conditions. The second group concerns variations of the state variables at the extremal points and defines the boundary optimality conditions for the adjoint variables.

By nullifying the coefficients of δt_0 and δt_f in the expression of $\delta \mathcal{J}^*$, one obtains [9]:

$$\begin{aligned}\frac{\partial \varphi}{\partial t_0} + \boldsymbol{\mu}^\top \frac{\partial \mathcal{X}}{\partial t_0} - \mathcal{H}_0 &= 0, \\ \frac{\partial \varphi}{\partial t_f} + \boldsymbol{\mu}^\top \frac{\partial \mathcal{X}}{\partial t_f} + \mathcal{H}_f &= 0.\end{aligned}$$

where \mathcal{H}_0 and \mathcal{H}_f denote the Hamiltonian evaluated at the initial and final instants, respectively.

These equations clarify the role of time constraints in the problem formulation. If time does not explicitly appear in the terminal function φ and is not constrained through \mathcal{X} , then the Hamiltonian must vanish at the corresponding boundary. In other words, for free initial or final time, the value of the Hamiltonian at that extremal point is determined by the optimization process and must satisfy the above conditions.

Conversely, if time is prescribed through explicit constraints (for example, $t_0 = a$ or $t_f = b$), then the corresponding Hamiltonian value is not fixed by a null condition, since the time variation is not admissible. Typical cases include fixed-duration transfers, where both t_0 and t_f are assigned, and problems with free final time, where only t_0 is fixed and t_f is subject to optimization.

By nullifying instead the coefficients of $\delta \mathbf{x}_0$ and $\delta \mathbf{x}_f$, the following conditions are obtained [9]:

$$\begin{aligned}\frac{\partial \varphi}{\partial \mathbf{x}_0} + \boldsymbol{\mu}^\top \frac{\partial \mathcal{X}}{\partial \mathbf{x}_0} + \boldsymbol{\lambda}_0^\top &= 0, \\ \frac{\partial \varphi}{\partial \mathbf{x}_f} + \boldsymbol{\mu}^\top \frac{\partial \mathcal{X}}{\partial \mathbf{x}_f} - \boldsymbol{\lambda}_f^\top &= 0.\end{aligned}$$

These expressions determine how the adjoint variables must be initialized and terminated at the extremal points. Their interpretation depends on whether a specific state component is fixed or free at the corresponding boundary.

If a given state variable x_i does not appear in the terminal function φ and is not constrained in \mathcal{X} at a certain boundary, then its associated adjoint variable must vanish at that point. On the other hand, if the state component is prescribed through boundary conditions, the corresponding adjoint variable remains free and is determined by the optimization process.

For instance, in a transfer problem where a particular state component is fixed at

both initial and final times, the associated adjoint variable is generally nonzero at both boundaries. If instead a state component is free at one boundary, the corresponding adjoint component must be zero at that point.

These relations provide a practical set of rules for assigning boundary conditions to the adjoint variables and for identifying whether the Hamiltonian must vanish at the extremal times. Together, the transversality and boundary optimality conditions complete the characterization of the solution at the trajectory endpoints and constitute an essential part of the TPBVP associated with the optimal control formulation.

4.1.3 Equations for adjoint and control variables

The differential relations governing the adjoint variables and the algebraic conditions defining the optimal control law follow from the nullification of the multiplying coefficients appearing in the last line of the first variation of the augmented merit function. Specifically, by imposing the vanishing of the coefficient associated with $\delta \mathbf{x}$, one obtains the set of Euler–Lagrange differential equations for the adjoint variables.

These equations describe how the adjoint vector evolves over time and read [9]:

$$\frac{d\boldsymbol{\lambda}}{dt} = -\left(\frac{\partial \mathcal{H}}{\partial \mathbf{x}}\right)^T.$$

The adjoint variables are therefore uniquely associated with the state variables, establishing a one-to-one correspondence between $\mathbf{x} \in \mathbb{R}^n$ and $\boldsymbol{\lambda} \in \mathbb{R}^n$. Together with the state equations, they form a coupled system of $2n$ first-order ordinary differential equations.

On the other hand, nullifying the coefficients multiplying $\delta \mathbf{u}$ yields a set of m algebraic relations for the control variables, namely [9]:

$$\left(\frac{\partial \mathcal{H}}{\partial \mathbf{u}}\right)^T = \mathbf{0},$$

which represent the stationarity condition of the Hamiltonian with respect to the control.

In the most general formulation, components of the control vector may be subject to admissibility constraints. In such a case, the admissible control set \mathcal{U} restricts the values that each control component may assume, typically through upper and lower bounds. The control may depend explicitly on time and implicitly on the state variables through the optimality conditions.

When explicit bounds are present, the optimal control must be selected so as to extremize the Hamiltonian within the admissible set. More precisely, according to Pontryagin’s Maximum Principle (PMP), for a maximization problem the optimal control $\mathbf{u}^*(t)$ satisfies [9]:

$$\mathcal{H}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \mathbf{u}^*, t) = \max_{\mathbf{u} \in \mathcal{U}} \mathcal{H}(\mathbf{x}^*, \boldsymbol{\lambda}^*, \mathbf{u}, t),$$

at every instant of time.

If the control is unconstrained and the Hamiltonian is regular with respect to \mathbf{u} , the extremum condition reduces to the stationarity relation:

$$\left(\frac{\partial \mathcal{H}}{\partial \mathbf{u}}\right)^T = \mathbf{0},$$

which determines the control algebraically. When admissibility bounds are imposed, the maximizer may instead lie at the boundary of the admissible set.

A particular situation arises when the Hamiltonian is affine with respect to a bounded control component u_i . In such a case [9]:

$$\frac{\partial \mathcal{H}}{\partial u_i} = \begin{cases} k_{u_i}, & \text{if } \mathcal{H} \text{ is affine with respect to } u_i, \\ f(u_i), & \text{otherwise,} \end{cases}$$

where k_{u_i} denotes a coefficient independent of u_i .

If the Hamiltonian is affine in u_i , the stationarity condition cannot determine the control uniquely (except in the degenerate case $k_{u_i} = 0$). The optimal strategy is then dictated by the sign of k_{u_i} . If $k_{u_i} > 0$, the Hamiltonian is maximized by selecting the upper admissible bound; if $k_{u_i} < 0$, the lower bound is selected. This switching behaviour corresponds to the so-called *bang–bang* control structure.

If instead $k_{u_i} = 0$ over a finite time interval, the problem enters a different regime in which the control cannot be determined from first-order conditions alone. Such intervals are referred to as singular arcs. In the present formulation, this situation does not arise and is therefore not further developed.

The set of relations obtained for transversality conditions, state optimality conditions, adjoint differential equations, and control algebraic equations collectively defines the structure of a TPBVP.

4.1.4 Multi-Point Boundary Value Problem Formulation

The general indirect OCP presented above is naturally introduced in a two-point boundary value form, since the trajectory is defined over a single time interval $[t_0, t_f]$ and the admissibility conditions are initially expressed only at the external boundaries.

However, when the control structure is prescribed *a priori* in terms of multiple thrust and coast arcs, the trajectory is no longer treated as a single uninterrupted segment. Instead, it is partitioned into a sequence of consecutive subintervals, each characterized by a specific control regime and separated by internal junction points.

In this case, the problem is more conveniently formulated as a *multi-point boundary value problem* (MPBVP). The trajectory is divided into n_p arcs, and the unknown solution must satisfy not only the boundary conditions at the initial and final times, but also a set of continuity and optimality conditions at the internal switching points. These additional conditions arise because the state and adjoint variables are propagated

arc-by-arc, and the complete solution must remain dynamically and variationally consistent across all arc junctions.

Let the j -th arc be defined over the interval [9]:

$$t \in [t_{(j-1)+}, t_{j-}], \quad j = 1, \dots, n_p,$$

where $t_{(j-1)+}$ and t_{j-} denote, respectively, the initial and final extremities of the j -th arc. The notation distinguishes the values immediately after and immediately before an internal boundary, so as to allow for the formal treatment of multipoint conditions at the switching instants. The corresponding state variables at the arc boundaries are denoted by:

$$\mathbf{x}_{(j-1)+}, \quad \mathbf{x}_{j-},$$

while the duration of each arc is:

$$\Delta t_j = t_{j-} - t_{(j-1)+}.$$

In general, these arc durations may be unknown and treated as optimization variables.

Once the trajectory is written arc-by-arc, the boundary conditions can no longer be expressed solely in terms of the two external endpoints. The complete set of constraints must include both the original endpoint conditions and the internal junction conditions between adjacent arcs. Accordingly, the boundary constraint vector takes the more general form [9]:

$$\mathcal{X}(\mathbf{x}_{(j-1)+}, \mathbf{x}_{j-}, t_{(j-1)+}, t_{j-}) = \mathbf{0}, \quad j = 1, \dots, n_p.$$

This notation emphasizes that, in the multipoint setting, the constraint function may depend on the state variables and time values at every arc boundary, not only at the global initial and final instants.

The merit index is correspondingly rewritten by summing the running contribution over all arcs. The Bolza functional therefore becomes [9]:

$$\mathcal{J} = \varphi(\mathbf{x}_{0-}, \mathbf{x}_{n_p+}, t_{0-}, t_{n_p+}) + \sum_{j=1}^{n_p} \int_{t_{(j-1)+}}^{t_{j-}} \Phi(\mathbf{x}(t), \mathbf{u}(t), t) dt.$$

The endpoint term retains the same conceptual role as in the TPBVP formulation, but the integral contribution is now explicitly decomposed arc-by-arc in order to account for the piecewise structure of the trajectory.

Introducing the adjoint variables and the vector of Lagrange multipliers associated with the complete set of multipoint constraints, the augmented merit functional becomes [9]:

$$\mathcal{J}^* = \varphi + \boldsymbol{\mu}^T \boldsymbol{\mathcal{X}} + \sum_{j=1}^{n_p} \int_{t_{(j-1)+}}^{t_{j-}} [\Phi + \boldsymbol{\lambda}^T (\mathbf{f} - \dot{\mathbf{x}})] dt.$$

As in the two-point case, integration by parts can be applied independently on each arc. This yields [9]:

$$\mathcal{J}^* = \varphi + \boldsymbol{\mu}^T \boldsymbol{\mathcal{X}} + \sum_{j=1}^{n_p} \left(\boldsymbol{\lambda}_{(j-1)+}^T \boldsymbol{x}_{(j-1)+} - \boldsymbol{\lambda}_{j-}^T \boldsymbol{x}_{j-} \right) + \sum_{j=1}^{n_p} \int_{t_{(j-1)+}}^{t_{j-}} \left(\boldsymbol{\Phi} + \boldsymbol{\lambda}^T \boldsymbol{f} - \dot{\boldsymbol{\lambda}}^T \boldsymbol{x} \right) dt.$$

The first-order variation is then obtained by considering the independent variations at every arc boundary. Proceeding exactly as in the TPBVP derivation, but now arc-by-arc, one obtains [9]:

$$\begin{aligned} \delta \mathcal{J}^* = & \left(\frac{\partial \varphi}{\partial t_{(j-1)+}} + \boldsymbol{\mu}^T \frac{\partial \boldsymbol{\mathcal{X}}}{\partial t_{(j-1)+}} - \mathcal{H}_{(j-1)+} \right) \delta t_{(j-1)+} + \\ & \left(\frac{\partial \varphi}{\partial t_{j-}} + \boldsymbol{\mu}^T \frac{\partial \boldsymbol{\mathcal{X}}}{\partial t_{j-}} + \mathcal{H}_{j-} \right) \delta t_{j-} + \\ & \left(\frac{\partial \varphi}{\partial \boldsymbol{x}_{(j-1)+}} + \boldsymbol{\mu}^T \frac{\partial \boldsymbol{\mathcal{X}}}{\partial \boldsymbol{x}_{(j-1)+}} + \boldsymbol{\lambda}_{(j-1)+}^T \right) \delta \boldsymbol{x}_{(j-1)+} + \\ & \left(\frac{\partial \varphi}{\partial \boldsymbol{x}_{j-}} + \boldsymbol{\mu}^T \frac{\partial \boldsymbol{\mathcal{X}}}{\partial \boldsymbol{x}_{j-}} - \boldsymbol{\lambda}_{j-}^T \right) \delta \boldsymbol{x}_{j-} + \\ & \sum_{j=1}^{n_p} \int_{t_{(j-1)+}}^{t_{j-}} \left[\left(\frac{\partial \mathcal{H}}{\partial \boldsymbol{x}} + \dot{\boldsymbol{\lambda}} \right)^T \delta \boldsymbol{x} + \frac{\partial \mathcal{H}}{\partial \boldsymbol{u}} \delta \boldsymbol{u} \right] dt. \end{aligned}$$

The stationarity of \mathcal{J}^* with respect to all independent variations yields the same differential conditions previously derived for the state, adjoint, and control variables, but now applied separately on each arc. Therefore, the Euler–Lagrange equations for the adjoints and the Hamiltonian stationarity condition retain exactly the same form as in the TPBVP formulation, with the only difference that they hold piecewise over each subinterval of the multipoint trajectory.

The main novelty of the MPBVP formulation lies in the treatment of the internal boundaries. At every junction between two adjacent arcs, one must impose continuity conditions for the state and, consistently with the variational structure, appropriate algebraic relations for the adjoint variables and Hamiltonian values. Expressing the conditions directly at the j -th internal boundary, one obtains [9]:

$$\begin{aligned} \frac{\partial \varphi}{\partial t_{j+}} + \boldsymbol{\mu}^T \frac{\partial \boldsymbol{\mathcal{X}}}{\partial t_{j+}} - \mathcal{H}_{j+} &= 0, & j = 0, \dots, n_p - 1, \\ \frac{\partial \varphi}{\partial t_{j-}} + \boldsymbol{\mu}^T \frac{\partial \boldsymbol{\mathcal{X}}}{\partial t_{j-}} + \mathcal{H}_{j-} &= 0, & j = 1, \dots, n_p, \\ \frac{\partial \varphi}{\partial \boldsymbol{x}_{j+}} + \boldsymbol{\mu}^T \frac{\partial \boldsymbol{\mathcal{X}}}{\partial \boldsymbol{x}_{j+}} + \boldsymbol{\lambda}_{j+}^T &= 0, & j = 0, \dots, n_p - 1, \\ \frac{\partial \varphi}{\partial \boldsymbol{x}_{j-}} + \boldsymbol{\mu}^T \frac{\partial \boldsymbol{\mathcal{X}}}{\partial \boldsymbol{x}_{j-}} - \boldsymbol{\lambda}_{j-}^T &= 0, & j = 1, \dots, n_p. \end{aligned}$$

These relations generalize the transversality and boundary optimality conditions to the multipoint setting. In practice, they provide the algebraic conditions that must

be satisfied at every arc junction, together with the continuity of the propagated state across adjacent arcs.

In the present work, this multipoint formulation is adopted whenever the switching structure is prescribed in advance. Each thrust or coast phase is treated as a distinct arc, integrated separately, and connected to the next one through internal continuity and switching conditions. The classical TPBVP therefore remains the natural general framework of the indirect OCP, whereas the implemented problem becomes an MPBVP whenever the trajectory is decomposed into multiple prescribed arcs. This formulation is particularly convenient for bang–bang problems, because it allows the switching instants to be treated explicitly as optimization variables while preserving the full analytical structure of the indirect method.

4.2 Spacecraft Trajectory Optimal Control Formulation

4.2.1 Adopted Dynamical Model and Problem Definition

The OCP introduced in the previous section is now formulated with explicit reference to the spacecraft dynamical model adopted in this work. The equations of motion governing the translational dynamics and mass evolution are those presented in Section 2, where the complete model including central gravity and perturbing accelerations has been defined.

The optimal control objective consists in determining the admissible thrust history that steers the spacecraft trajectory while maximizing the final mass at the terminal time. In accordance with the Mayer formulation discussed previously, no running cost is introduced and the performance index reduces to a pure endpoint contribution. The merit index therefore reads:

$$\mathcal{J} = m(t_f),$$

which expresses directly the maximization of the remaining propellant at the end of the maneuver.

The control variable corresponds to the thrust vector, whose magnitude and direction are bounded by physical limitations of the propulsion system. The control law must satisfy the spacecraft equations of motion while fulfilling the boundary conditions imposed on the trajectory.

The state vector is expressed in spherical ZEN coordinates, consistently with the adopted dynamical framework. It is defined as:

$$\mathbf{x} = [r \ \vartheta \ \varphi \ u \ v \ w \ m]^T,$$

where the first three components describe the spacecraft position, the following three represent the velocity components resolved in the ZEN reference frame, and the last component corresponds to the spacecraft mass.

Each state variable is associated with a corresponding adjoint variable, introduced

through the variational formulation of the OCP. The augmented state vector is therefore written as:

$$\mathbf{y} = [r \ \vartheta \ \varphi \ u \ v \ w \ m \ \lambda_r \ \lambda_\vartheta \ \lambda_\varphi \ \lambda_u \ \lambda_v \ \lambda_w \ \lambda_m]^T.$$

At the continuous level, the problem is therefore defined by the coupled system of state and adjoint differential equations, together with the boundary and optimality conditions derived in the previous section. For a single-arc trajectory this structure corresponds directly to a TPBVP. When instead the thrust profile is represented through a prescribed sequence of thrust and coast arcs, the same canonical dynamics is enforced piecewise over the successive subintervals and the implemented problem takes the form of a MPBVP, with internal junction conditions imposed at the arc boundaries.

The solution provides simultaneously the optimal trajectory and the associated control history that maximizes the final spacecraft mass.

4.2.2 Hamiltonian and Switching Function

By explicitly substituting the spacecraft equations of motion into the general Hamiltonian definition $\mathcal{H} = \lambda^T f(\mathbf{x}, \mathbf{u}, t)$, the Hamiltonian associated with the present dynamical model can be written in expanded form. The derivation follows directly from the variational formulation introduced in the previous section and applies to the canonical state-adjoint system on each arc of the trajectory.

Recalling that the state is expressed in spherical ZEN coordinates and that the velocity components are (u, v, w) , the Hamiltonian becomes [9]:

$$\begin{aligned} \mathcal{H} = & \lambda_r u + \lambda_\vartheta \frac{v}{r \cos \varphi} + \lambda_\varphi \frac{w}{r} \\ & + \lambda_u \left(-\frac{\mu}{r^2} + \frac{v^2}{r} + \frac{w^2}{r} + \frac{T_u}{m} + (a_p)_u \right) \\ & + \lambda_v \left(-\frac{uv}{r} - \frac{vw}{r} \tan \varphi + \frac{T_v}{m} + (a_p)_v \right) \\ & + \lambda_w \left(-\frac{uw}{r} + \frac{v^2}{r} \tan \varphi + \frac{T_w}{m} + (a_p)_w \right) - \lambda_m \frac{T}{c}. \end{aligned}$$

This expression explicitly shows the contribution of kinematic terms, gravitational acceleration, perturbing accelerations, thrust components, and mass variation. Each adjoint variable multiplies the corresponding state equation, preserving the canonical structure derived in the general optimal control formulation.

For compactness, the Hamiltonian may be reorganized by separating the purely dynamical terms from those depending on the thrust-to-mass ratio. Introducing the velocity adjoint vector:

$$\lambda_V = [\lambda_u \ \lambda_v \ \lambda_w]^T,$$

the Hamiltonian can be rewritten as [9]:

$$\mathcal{H} = \lambda_r^T \mathbf{V} + \lambda_V^T \left(-\frac{\mu}{r^3} \mathbf{r} + \mathbf{a}_p \right) + \frac{T}{m} \mathcal{S}_{\mathcal{F}},$$

where all the terms multiplying the thrust magnitude have been grouped into the switching function [9]:

$$\mathcal{S}_{\mathcal{F}} = \lambda_V^T \hat{\mathbf{T}} - \lambda_m \frac{m}{c}.$$

The control vector is composed of the thrust magnitude T and its direction $\hat{\mathbf{T}}$. Since the Hamiltonian is linear with respect to the thrust magnitude, the Pontryagin Maximum Principle implies that the optimal thrust level is determined by the sign of the switching function. In particular, the optimal thrust magnitude must maximize the Hamiltonian within the admissible set $0 \leq T \leq T_{\max}$. This leads to the bang-bang control law:

$$T = \begin{cases} T_{\max}, & \mathcal{S}_{\mathcal{F}} > 0, \\ 0, & \mathcal{S}_{\mathcal{F}} < 0, \end{cases}$$

which corresponds to alternating thrust and coast arcs along the trajectory. Singular arcs, associated with the case $\mathcal{S}_{\mathcal{F}} = 0$ over a finite interval, are not considered in the present formulation.

Concerning the thrust direction, maximization of the Hamiltonian requires maximizing the scalar product $\lambda_V^T \hat{\mathbf{T}}$. This condition is satisfied when the thrust unit vector is aligned with the adjoint velocity vector. Defining the primer vector as:

$$\lambda_V = \begin{bmatrix} \lambda_u \\ \lambda_v \\ \lambda_w \end{bmatrix}, \quad \lambda_V = \|\lambda_V\|,$$

the optimal thrust direction is therefore parallel to λ_V . With this alignment, the switching function assumes the scalar form [9]:

$$\mathcal{S}_{\mathcal{F}} = \lambda_V - \lambda_m \frac{m}{c}.$$

The switching function thus governs the activation or deactivation of thrust along the trajectory and plays a central role in determining the sequence of thrust and coast phases characterizing the optimal solution.

4.2.3 Optimal Thrust Direction and Angular Parametrization

The thrust vector is defined by its magnitude T and by its direction in the ZEN reference frame. To provide a convenient parametrization of its orientation, two angles are introduced: the thrust elevation angle α_T and the heading angle β_T .

Figure 4.1 illustrates the geometric meaning of these angles and the decomposition of the thrust vector along the ZEN basis directions $(\hat{u}, \hat{v}, \hat{w})$.

With this parametrization, the thrust vector components can be written as [9]:

$$\mathbf{T} = \begin{bmatrix} T_u \\ T_v \\ T_w \end{bmatrix} = T \begin{bmatrix} \sin \alpha_T \\ \cos \alpha_T \cos \beta_T \\ \cos \alpha_T \sin \beta_T \end{bmatrix}, \quad T = \|\mathbf{T}\|. \quad (4.4)$$

The optimal thrust direction follows from the maximization of the Hamiltonian with respect to the control direction. From the Hamiltonian structure derived in the previous




Figure 4.1: Thrust orientation in the ZEN reference frame and definition of elevation angle α_T and heading angle β_T .

section, this requirement reduces to maximizing the scalar product between the adjoint velocity vector and the thrust unit vector. Consequently, the thrust direction must be aligned with the adjoint velocity vector, also known as the primer vector.

Let the adjoint velocity vector be defined as [9]:

$$\lambda_V = \begin{bmatrix} \lambda_u \\ \lambda_v \\ \lambda_w \end{bmatrix}, \quad \lambda_V = \|\lambda_V\|.$$

Under the optimality condition, the thrust unit vector coincides with the normalized primer vector. By equating the thrust direction components in (4.4) with the normalized adjoint vector, the following relations are obtained [9]:

$$\begin{aligned} \sin \alpha_T &= \frac{\lambda_u}{\lambda_V}, \\ \cos \alpha_T \cos \beta_T &= \frac{\lambda_v}{\lambda_V}, \\ \cos \alpha_T \sin \beta_T &= \frac{\lambda_w}{\lambda_V}, \end{aligned} \tag{4.5}$$

where $\lambda_V = \sqrt{\lambda_u^2 + \lambda_v^2 + \lambda_w^2}$ denotes the magnitude of the adjoint velocity vector.

These relations show that the thrust direction coincides with the normalized adjoint velocity vector. Therefore, the angular parametrization is fully consistent with the alignment condition between the thrust unit vector and the primer vector derived from the Pontryagin Maximum Principle.

In practice, once the adjoint variables are known at a given instant, the optimal thrust orientation is determined algebraically through relations (4.5). The time evolution of the thrust direction therefore follows implicitly from the propagation of the adjoint variables governed by the adjoint differential equations introduced earlier. In the presence of a multi-arc trajectory structure, these relations apply independently on each arc of the multipoint boundary value problem.

4.2.4 Switching Structure and Imposed Arc Strategy

Although the sign of the switching function $\mathcal{S}_{\mathcal{F}}$ determines the optimal thrust level according to the PMP, its numerical evaluation may become delicate in high-fidelity dynamical environments. In strongly nonlinear regimes, small variations in the state or adjoint variables can induce rapid oscillations of $\mathcal{S}_{\mathcal{F}}$ around zero, leading to undesired thrust insertions or eliminations during numerical integration.

To improve robustness and convergence properties, the thrust-coast structure is prescribed *a priori*. The trajectory is partitioned into a sequence of thrust and coast arcs, whose number is selected based on physical insight into the maneuver. Each arc has an associated duration Δt_j , treated as an optimization variable.

Typical structures may include, for instance:

- single-burn: thrust \rightarrow coast;
- multi-burn: alternating arcs with a prescribed number of switches.

Within this framework, the canonical state-adjoint system is integrated separately over each arc. Continuity of the state variables is enforced at the internal arc boundaries, ensuring that the overall trajectory remains dynamically consistent across successive segments. Additional conditions are imposed at the switching times by requiring the vanishing of the switching function, which guarantees consistency with the Pontryagin Maximum Principle at the transition between thrust and coast phases.

As a consequence, the boundary-value problem no longer involves only the initial and final conditions but also includes internal junction conditions at the switching times. The resulting formulation therefore takes the form of a MPBVP, in which the arc durations and the boundary conditions at all junction points are solved simultaneously.

The imposed switching structure significantly enhances numerical stability, particularly in scenarios involving complex gravitational perturbations. If the prescribed sequence of arcs satisfies the PMP optimality conditions, the resulting trajectory corresponds to the optimal solution. Otherwise, the behaviour of the switching function provides guidance for modifying the arc structure and refining the solution through additional switching points or different arc configurations.

This strategy reduces the sensitivity of the solution to local numerical fluctuations of the switching function and improves the reliability of indirect optimization when high-fidelity dynamical models are employed.

4.2.5 Adjoint Equations

To complete the optimal control formulation, the explicit expressions of the adjoint dynamics associated with the adopted spherical ZEN state are reported. These equations are obtained by applying the Euler–Lagrange relation to the Hamiltonian of the problem and define the differential system that must be integrated together with the state

equations under the optimal control law. The adjoint dynamics follows from [9]:

$$\dot{\lambda} = -\left(\frac{\partial \mathcal{H}}{\partial \mathbf{x}}\right)^\top,$$

and, for the spherical ZEN model introduced in Section 2.2, assumes the following explicit form [9]:

$$\begin{aligned} \dot{\lambda}_r = \frac{1}{r^2} & \left[\lambda_\vartheta \frac{v}{\cos \varphi} + \lambda_\varphi w + \lambda_u \left(-\frac{2\mu}{r} + v^2 + w^2 \right) + \lambda_v (-uv + vw \tan \varphi) + \lambda_w (-uw - v^2 \tan \varphi) \right] \\ & - \lambda_u \frac{\partial(a_p)_u}{\partial r} - \lambda_v \frac{\partial(a_p)_v}{\partial r} - \lambda_w \frac{\partial(a_p)_w}{\partial r} - \mathcal{S}_{\mathcal{F}} \frac{\partial T}{\partial r}, \end{aligned} \quad (4.6)$$

$$\dot{\lambda}_\vartheta = -\lambda_u \frac{\partial(a_p)_u}{\partial \vartheta} - \lambda_v \frac{\partial(a_p)_v}{\partial \vartheta} - \lambda_w \frac{\partial(a_p)_w}{\partial \vartheta} + \mathcal{S}_{\mathcal{F}} \frac{\partial T}{\partial \vartheta}, \quad (4.7)$$

$$\begin{aligned} \dot{\lambda}_\varphi = \frac{1}{r \cos^2 \varphi} & \left(-\lambda_\vartheta v \sin \varphi - \lambda_v v w + \lambda_w v^2 \right) \\ & - \lambda_u \frac{\partial(a_p)_u}{\partial \varphi} - \lambda_v \frac{\partial(a_p)_v}{\partial \varphi} - \lambda_w \frac{\partial(a_p)_w}{\partial \varphi} + \mathcal{S}_{\mathcal{F}} \frac{\partial T}{\partial \varphi}, \end{aligned} \quad (4.8)$$

$$\dot{\lambda}_u = \frac{1}{r} (-\lambda_r r + \lambda_v v + \lambda_w w),$$

$$\dot{\lambda}_v = \frac{1}{r} \left[-\frac{\lambda_\vartheta}{\cos \varphi} - 2\lambda_u v + \lambda_v (u - w \tan \varphi) + 2\lambda_w v \tan \varphi \right],$$

$$\dot{\lambda}_w = \frac{1}{r} (-\lambda_\varphi - 2\lambda_u w - \lambda_v v \tan \varphi + \lambda_w u),$$

$$\dot{\lambda}_m = \frac{T}{m^2} \lambda_v - \lambda_u \frac{\partial(a_p)_u}{\partial m} - \lambda_v \frac{\partial(a_p)_v}{\partial m} - \lambda_w \frac{\partial(a_p)_w}{\partial m}. \quad (4.9)$$

The above expressions explicitly account for the dependence of the perturbing accelerations and of the thrust magnitude on the state variables. In the present formulation, the thrust magnitude is treated as an explicit control variable and does not depend on the state; therefore:

$$\frac{\partial T}{\partial r} = \frac{\partial T}{\partial \vartheta} = \frac{\partial T}{\partial \varphi} = 0,$$

and the corresponding terms in (4.6), (4.7), and (4.8) vanish. Similarly, if the selected perturbing accelerations are independent of the spacecraft mass, then $\partial(a_p)_\bullet / \partial m = 0$, and equation (4.9) reduces to:

$$\dot{\lambda}_m = \frac{T}{m^2} \lambda_v.$$

The adjoint system, together with the state equations and the optimal control law derived from the switching function, defines the canonical dynamical system governing the optimal trajectory. In the presence of a prescribed multi-arc thrust structure, these equations are integrated independently on each arc of the trajectory, while continuity conditions are enforced at the switching times within the multipoint boundary value problem formulation.

4.2.6 Numerical Solution via Shooting and Differential Correction

The boundary value problem arising from the indirect optimal control formulation does not require the determination of arbitrary functions, but rather the identification of a finite set of boundary parameters that uniquely determine the state-costate evolution satisfying all optimality and boundary conditions. In indirect approaches, this finite-dimensional nonlinear problem is typically solved through a shooting procedure.

In the shooting framework, the canonical state-adjoint system is integrated forward in time starting from an initial guess of the unknown boundary parameters. The resulting trajectory generally does not satisfy the terminal constraints and transversality conditions; therefore, an iterative correction of the initial guess is required. This process converts the boundary value problem — whether in two-point or multi-point form — into a nonlinear root-finding problem.

A design vector \mathbf{p} is introduced to collect all unknown boundary quantities. In the present formulation, it may include the initial adjoint variables and the durations or locations of switching times, namely [10]:

$$\mathbf{p} = [\lambda(t_0)^\top, t_1, \dots, t_{N_s}, t_f]^\top,$$

where t_1, \dots, t_{N_s} denote the imposed switching times associated with the selected thrust-coast structure, and t_f is included if the final time is free.

For a given guess \mathbf{p} , the canonical system is propagated using the closed-loop optimal control law derived from the PMP. When a multi-arc thrust structure is prescribed, the propagation is performed arc by arc over the intervals defined by the switching times. At the end of the propagation, a residual vector $F(\mathbf{p})$ is assembled. This vector collects all boundary mismatches, including:

- terminal state constraints $\mathcal{X} = \mathbf{0}$;
- transversality conditions ($\mathcal{H}(t_f) = 0$ when t_f is free, and the prescribed value of $\lambda_m(t_f)$ for the Mayer problem);
- continuity conditions of the state variables at the switching times;
- switching conditions enforcing the vanishing of the switching function at each switching time:

$$\mathcal{S}_{\mathcal{F}}(t_i) = 0, \quad i = 1, \dots, N_s.$$

The numerical problem then reduces to finding a vector \mathbf{p}^* such that [10]:

$$F(\mathbf{p}^*) = \mathbf{0}.$$

Implementation in OCULUS

In the present work, the differential correction procedure is implemented using the OCULUS solver [10]. Within this framework, the user specifies the problem definition, including the design vector components and boundary conditions, while the solver automatically:

- integrates the canonical system (arc by arc when a multi-arc structure is prescribed);
- assembles the residual vector F ;
- computes or updates the Jacobian J ;
- applies a suitable correction rule to obtain the updated guess.

This structure allows different correction strategies and regularization modes to be employed without altering the optimal control formulation itself. As a result, the indirect shooting method can be applied to high-fidelity dynamical models while maintaining numerical robustness and flexibility in the treatment of switching times and boundary conditions.

5 Optimal Solution in the Unperturbed Dynamical Model

This chapter presents the optimal solution of the collision-avoidance problem under the unperturbed dynamical model. The objective is twofold: first, to document how the indirect OCP formulated in the previous chapter is concretely implemented within the OCULUS environment by specifying the state, adjoint variables, boundary conditions, and switching structure; second, to provide a clear physical interpretation of the resulting trajectory, control history, and adjoint behavior when the spacecraft motion is governed solely by the two-body gravitational field.

The choice of starting from the unperturbed two-body model is deliberate. In the absence of additional perturbations, the dynamical structure is simpler and the interaction between state and adjoint dynamics becomes more transparent. This setting allows the switching structure, thrust-direction behavior, and adjoint evolution to be analyzed without secondary effects that could obscure the fundamental mechanisms governing the optimal solution. Moreover, the unperturbed case provides a consistent reference configuration that will later serve as a benchmark for the perturbed dynamical model.

From a numerical standpoint, the two-body formulation also exhibits significantly more regular convergence properties than the fully perturbed problem. The associated boundary-value problem, which in this case reduces to a relatively simple multi-point formulation, is less sensitive to the initial costate guess and admits a more stable differential-correction process. This facilitates the identification of a consistent initial adjoint set and provides an informed estimate of the magnitude and structure of the costate variables. Such information is essential when progressively introducing additional dynamical effects, since attempting to identify a suitable initial guess directly in the perturbed model would require exploring a substantially larger and less structured parameter space.

The development follows a progressive continuation process. The analysis starts from a deliberately simple baseline configuration, namely a circular equatorial orbit in a free-final-time formulation. This choice leads to a well-conditioned boundary-value problem and allows the identification of a consistent initial costate set. Starting from this reference solution, the problem is gradually continued toward a fully three-dimensional configuration by introducing orbital inclination, while preserving the same optimal control structure.

Once the inclined reference case is converged, the terminal constraint is updated to enforce the planar displacement selected in the collision-avoidance analysis within

the encounter plane. The formulation is still treated with free final time, so that the maneuver duration naturally emerges from the optimality conditions rather than being imposed a priori.

The second part of the chapter addresses the operational requirement of achieving the prescribed displacement at a specific epoch. The OCP is therefore reformulated with fixed final time aligned with the collision time. In this configuration, the maneuver start time becomes an optimization variable and the corresponding transversality condition at the initial boundary is enforced. The resulting fixed-time solution is then analyzed in detail through state, control, and costate histories, including a discussion of thrust-direction reconstruction from the primer vector and an angular-consistency verification at the terminal time.

Finally, the same unperturbed setup is used to investigate how the optimal solution changes when the propulsion parameters are modified to represent electric-propulsion regimes, while keeping the dynamical model and boundary conditions unchanged.

5.1 Implementation within the OCULUS Solver

The OCP is solved using the OCULUS indirect solver [10], which is designed to handle boundary-value problems derived from Pontryagin’s Maximum Principle, including both two-point and multi-point formulations. Within this framework, the canonical state-costate system is integrated and a finite set of boundary parameters is iteratively adjusted until all optimality and endpoint conditions are simultaneously satisfied.

From an implementation standpoint, the problem must be fully specified before being passed to the solver. In particular, OCULUS requires a complete definition of the dynamical model, the Hamiltonian structure, the boundary constraints, and the thrust-arc configuration. The following elements are therefore provided:

- the initial spacecraft state $x(t_0)$, defining the starting point of the trajectory in the unperturbed two-body model;
- the full set of state differential equations, corresponding to the spherical ZEN formulation under the central gravitational field;
- the complete set of adjoint differential equations, obtained from the Hamiltonian through the Euler–Lagrange relations;
- the Hamiltonian definition and the associated switching function, which determine the optimal thrust magnitude and direction along each arc;
- the boundary conditions, including terminal constraints, transversality conditions, and, when applicable, switching conditions;
- the propulsion parameters, namely the maximum thrust T_{\max} and exhaust velocity c , which define the control authority and the propellant mass-depletion law;
- the selected switching structure, specified in terms of the number of thrust and coast arcs and the corresponding switching times.

With these elements defined, the OCP is fully characterized at the continuous level. OCULUS then treats the unknown initial costates and, when applicable, the arc durations or switching times as free parameters, adjusting them until the canonical system satisfies all boundary and optimality conditions. In practice, the canonical equations are integrated arc by arc according to the prescribed thrust-coast structure, while the solver iteratively corrects the boundary parameters until the residual vector associated with the boundary conditions vanishes. In this sense, the solver does not introduce additional modeling assumptions, but operates directly on the analytical structure derived in the previous chapter.

5.2 Propulsion Model and Parameter Selection

5.2.1 Baseline Propulsion Configuration

The spacecraft considered in this study has an initial mass:

$$m_0 = 462 \text{ kg.}$$

This mass level places the vehicle in the class of compact microsatellites or small satellites. For a spacecraft of this size, collision-avoidance maneuvers must provide sufficient acceleration authority to generate a measurable orbital displacement within operationally realistic time windows around the time of closest approach. In practical conjunction scenarios, maneuver lead times are typically limited to a few orbital periods, and the propulsion system must therefore deliver a finite velocity increment over time intervals that remain compatible with onboard operational constraints.

For the baseline unperturbed analysis, a compact chemical propulsion system representative of mono or bipropellant technology is adopted. According to the NASA small spacecraft propulsion state-of-the-art survey [28], systems in this class typically operate in the sub-newton to few-newton thrust range and achieve specific impulse values on the order of 300–320 s. These characteristics make them particularly suitable for short-duration orbit-adjustment or collision-avoidance maneuvers, where a relatively concentrated thrust arc is preferable to prolonged low-acceleration thrusting.

Within this technological envelope, the propulsion parameters are selected as:

$$T_{\max} = 0.5 \text{ N}, \quad c = 3000 \text{ m/s,}$$

corresponding to:

$$I_{sp} = \frac{c}{g_0} \approx 306 \text{ s.}$$

The selected thrust level of 0.5 N falls within the sub-newton regime typically associated with compact chemical propulsion systems for small satellites [28]. For the assumed initial mass, this corresponds to a maximum acceleration:

$$a_{\max} = \frac{T_{\max}}{m_0} \approx 1.08 \times 10^{-3} \text{ m/s}^2.$$

Although this acceleration is several orders of magnitude smaller than the central gravitational acceleration along the orbit, it is sufficiently large to produce velocity

increments of several tenths of meters per second over time intervals of a few hundred seconds. Such velocity changes are entirely consistent with the scale of orbital displacements required for probabilistic collision mitigation in low Earth orbit. At the same time, the thrust authority remains bounded and realistic, ensuring that the optimal control solution does not rely on artificially high accelerations that would distort the switching structure or compress the maneuver into unrealistically short time spans.

The exhaust velocity $c = 3000$ m/s, corresponding to a specific impulse of approximately 306 s, lies within the upper performance range of high-efficiency chemical monopropellant or bipropellant systems for small spacecraft. This value provides a realistic compromise between thrust capability and propellant efficiency. Since the adopted performance index maximizes final mass, the exhaust velocity directly influences the mass depletion rate and therefore affects the adjoint dynamics associated with the mass state. Consequently, the selected value of c has a direct impact on the structure of the switching function and on the allocation of thrust within the optimal solution.

The adoption of this chemical propulsion configuration for the initial solution search is motivated by numerical and structural considerations. A thrust level that is moderate relative to the spacecraft mass leads to a burn duration that is short compared to the characteristic orbital timescale. As a result, the maneuver can be interpreted as a nearly impulsive energy injection followed by a long natural drift under Keplerian dynamics. This separation of timescales clarifies the physical mechanism underlying the optimal solution and makes the role of the adjoint variables more transparent. From a numerical standpoint, the short thrust phase limits the time interval over which the control actively reshapes the trajectory. Consequently, deviations in the initial adjoint guess have less time to accumulate during the controlled portion of the motion, which generally improves the robustness of the shooting procedure. The resulting configuration therefore provides a stable reference case against which alternative propulsion regimes, such as lower-thrust electric systems, can later be evaluated while keeping the dynamical model unchanged.

5.2.2 Continuation toward Electric Propulsion Regimes

Once a first converged solution is obtained with the chemical baseline configuration, continuation is performed by progressively modifying the propulsion parameters in order to explore different propulsion regimes while preserving the same dynamical framework and boundary conditions. This approach allows the influence of propulsion characteristics to be isolated, since the gravitational model, state formulation, and optimal control structure remain unchanged.

Two additional propulsion configurations representative of electric propulsion technologies are considered. The first corresponds to a Hall effect thruster regime, characterized by a moderate thrust level and significantly higher exhaust velocity, with parameters:

$$T_{\max} = 0.05 \text{ N}, \quad c = 20000 \text{ m/s.}$$

The second represents a gridded ion thruster configuration, operating at even lower thrust and higher specific impulse, with:

$$T_{\max} = 0.01 \text{ N}, \quad c = 30000 \text{ m/s.}$$

These values are consistent with the performance ranges reported in the NASA small spacecraft propulsion survey [28] and reflect the transition from moderate-thrust chemical propulsion to strongly low-thrust electric systems.

For the 462 kg spacecraft considered here, the corresponding thrust-to-mass ratios decrease by one to two orders of magnitude relative to the chemical baseline. Consequently, the characteristic time required to generate a given velocity increment increases proportionally. While the chemical configuration can produce sub-meter-per-second velocity changes over relatively short time intervals, the electric configurations require substantially longer thrust durations to achieve comparable orbital energy variations.

From an optimal control perspective, reducing the available thrust primarily affects the temporal distribution of the control. The fuel-optimal problem retains its bang–bang structure, since the control remains bounded and the switching condition $\mathcal{S}_{\mathcal{F}} = 0$ is unchanged. However, as the thrust-to-mass ratio decreases, the required energy increment must be accumulated progressively rather than through a short concentrated burn. The maneuver therefore transitions from an impulsive-like strategy toward an extended low-thrust regime in which orbital reshaping occurs over a significant fraction of the trajectory.

The increase in specific impulse modifies the role of the mass dynamics within the Hamiltonian. Since the mass flow rate is inversely proportional to the exhaust velocity, a higher I_{sp} reduces the instantaneous propellant consumption associated with thrusting. This changes the balance encoded in the switching function between acceleration gain and mass expenditure, thereby affecting the adjoint variable λ_m and its contribution to the optimal control law.

Although the two-body dynamical model remains unchanged, the indirect numerical solution typically becomes more sensitive as thrust decreases. This increased sensitivity arises from the longer integration intervals over which the control remains active, so that small variations in the initial adjoint guess propagate over extended time spans and can significantly affect the terminal residuals.

The comparison among propulsion regimes therefore provides a structured way to assess how the optimal collision-avoidance strategy evolves as the propulsion system transitions from impulsive-like to strongly low-thrust behavior. The detailed numerical results for each configuration are reported in the following sections.

5.3 Baseline Circular-Orbit Configuration

This section presents the first converged optimal solution obtained in the unperturbed dynamical model. This configuration is adopted as the *baseline reference* for the entire continuation process developed in the following sections.

The initial case is intentionally selected as a circular equatorial orbit. The combined assumptions of zero eccentricity and zero inclination lead to a dynamically symmetric configuration in spherical ZEN coordinates. In this configuration, the absence of perturbations and the equatorial symmetry reduce the number of active dynamical

mechanisms involved in satisfying the terminal constraint. The optimal maneuver is essentially driven by an energy variation mechanism, primarily associated with the tangential velocity component. This structural simplicity makes the qualitative role of the different costate components more transparent. Such transparency is particularly advantageous in an indirect shooting formulation, where the initial adjoint guess must be inferred without prior knowledge of its magnitude or relative structure.

A further fundamental motivation for starting from this simplified configuration lies in the intrinsic difficulty of determining appropriate initial guesses for the costates. In an indirect formulation, the initial adjoint vector λ_0 is unknown and must be identified within a high-dimensional solution space. Without prior information on their order of magnitude or relative structure, the shooting method may explore a very wide region of the parameter space, leading to poor conditioning or divergence. By first solving a dynamically simple and symmetric problem, it becomes possible to identify the typical scale, sign, and relative dominance of the different costate components. This provides valuable insight into the structure of the optimal solution and supplies a physically meaningful initial guess for more complex configurations.

Starting from this baseline therefore serves a dual purpose: it yields a numerically robust solution and, at the same time, establishes a reference adjoint structure that can be progressively continued toward inclined or perturbed cases. This strategy significantly improves convergence speed and robustness in the subsequent continuation process. In addition, the baseline case provides a physically transparent reference to interpret the structure of the optimal policy (thrust direction, switching behaviour, and adjoint trends) before moving to more demanding configurations.

5.3.1 Free Final Time Formulation: Problem Definition and Switching Structure

The reference orbit is defined as a circular equatorial orbit at 600 km altitude. In spherical ZEN coordinates, the initial condition is:

$$\mathbf{x}_0 = [r_0, \vartheta_0, \varphi_0, u_0, v_0, w_0, m_0]^T = [R_E + 600 \text{ km}, 0, 0, 0, \sqrt{\mu/r_0} \text{ m/s}, 0, 462.0 \text{ kg}]^T.$$

The condition $u_0 = 0$ ensures zero radial velocity, while the tangential velocity $v_0 = \sqrt{\mu/r_0}$ guarantees circular motion in the two-body model. The choice $w_0 = 0$ indicates that, at the initial time, the velocity has no component along the φ direction of the ZEN frame. In the present equatorial configuration, the φ direction is normal to the orbital plane; therefore, $w_0 = 0$ implies that the motion is initially confined to the orbital plane. It should be emphasized that w is not intrinsically an out-of-plane velocity component in the general ZEN formulation; this interpretation holds here only because the reference orbit is equatorial.

For completeness, the equivalent Keplerian elements are:

$$[a, e, i, \Omega, \omega, \nu] = [r_0, 0, 0, 0, 0, 0].$$

The terminal requirement imposed in this first problem consists of a small radial increase:

$$r_f = r_0 + 100 \text{ m}.$$

All other terminal state components are left free.

In an unperturbed two-body setting, a circular orbit satisfies:

$$\mathcal{E} = -\frac{\mu}{2a},$$

where \mathcal{E} is the specific mechanical energy and a is the semi-major axis. For near-circular motion, a small increase in orbital radius corresponds approximately to an increase in the semi-major axis and therefore in the specific orbital energy. From a physical standpoint, this energy increment is most efficiently produced by a short thrust arc with a dominant tangential component, since a tangential acceleration directly increases orbital energy while avoiding unnecessary excitation of radial oscillations. Once the energy increment has been injected, the spacecraft can coast freely until the radial coordinate reaches the prescribed value at some final time.

Because only $r(t_f)$ is constrained, the boundary-value problem does not require a terminal thrust arc to “shape” the endpoint state: a single thrust arc followed by a coast arc is sufficient to satisfy the terminal condition, and this motivates the adoption of a minimal switching structure.

The final time is treated as a free variable. As a consequence, the transversality condition for free t_f applies. For a Mayer fuel-optimal problem, this yields:

$$\mathcal{H}(t_f) = 0, \tag{5.1}$$

which must be satisfied together with the terminal radius constraint and the remaining endpoint conditions of the problem. Since both the switching time t_1 and the final time t_f are unknown, these additional scalar conditions are required to close the resulting multi-point boundary-value problem consistently within the indirect formulation.

A thrust-coast switching structure is imposed:

$$T(t) = \begin{cases} T_{\max}, & t_0 \leq t \leq t_1, \\ 0, & t_1 < t \leq t_f, \end{cases}$$

where the switching time t_1 is an unknown parameter of the problem. The optimality condition associated with bang–bang control requires the switching function to vanish at the switching instant:

$$\mathcal{S}_{\mathcal{F}}(t_1) = 0. \tag{5.2}$$

Equation (5.2) represents an additional scalar condition that must be satisfied together with the terminal constraint and the free-final-time condition (5.1). It ensures that the control transition is optimal and is explicitly enforced within the shooting residual vector.

Overall, the baseline problem is therefore characterized by a short energy-injection phase followed by a long natural drift under Keplerian dynamics, until the radial condition is met at the final time.

5.3.2 Converged Solution

Let the initial adjoint vector be defined as:

$$\lambda_0 = \left[\lambda_r, \lambda_\vartheta, \lambda_\varphi, \lambda_u, \lambda_v, \lambda_w, \lambda_m \right]_{t=t_0}^T.$$

The converged baseline solution yields:

$$\lambda_0 = \left[0.000125099, 0, 0, -0.000521659, 0.1539973, 0, 0.999990974 \right]^T.$$

This solution is numerically stable and is adopted as the reference point for the continuation strategy developed in the following sections.

The structure of this adjoint vector is consistent with both the geometry of the baseline configuration and the physical nature of the maneuver. The dominant positive value of $\lambda_v(t_0)$ indicates that variations in the tangential velocity component have the strongest impact on satisfying the terminal radial constraint at minimum propellant cost. This is coherent with the structure of the optimal burn: in order to increase the orbital radius in a near-circular two-body orbit, the most efficient action is to inject energy through a predominantly tangential acceleration. Accordingly, the tangential state component v plays the primary role in shaping the terminal radial condition, and its associated costate is dominant.

The small but nonzero value of $\lambda_u(t_0)$ is also physically meaningful. Variations in the radial velocity component do not directly modify the orbital energy in the same way as tangential velocity variations; however, they affect the phase of the subsequent Keplerian motion and therefore influence the timing at which the radial extremum is reached. Since the final time is free and the terminal constraint is imposed on $r(t_f)$, the initial radial velocity contributes indirectly to the satisfaction of the endpoint condition. This explains why λ_u is nonzero, although significantly smaller than λ_v .

Conversely, $\lambda_w(t_0) = 0$ reflects the fact that the optimal solution does not exploit the third velocity component of the ZEN frame. In the general ZEN formulation, w is the velocity component associated with the φ direction. In the present baseline configuration, since the orbit is equatorial and remains confined to a fixed plane in the two-body model, the φ direction corresponds with the direction normal to the orbital plane. Therefore, in this specific case, w effectively represents an out-of-plane velocity component. Because the terminal requirement concerns only a radial increase and does not involve any change in orbital plane, any nonzero w would introduce a deviation from planarity without contributing to the satisfaction of the radial constraint. Such a deviation would only increase propellant consumption. The optimal solution consequently remains planar in this equatorial setting, and the associated costate λ_w converges to zero.

The null values of $\lambda_\vartheta(t_0)$ and $\lambda_\varphi(t_0)$ must be interpreted in terms of the sensitivity of the optimal cost with respect to the corresponding state variables. In the present Mayer formulation, the performance index depends only on the final mass, and the only terminal constraint is imposed on the radial coordinate r_f . The costates λ_ϑ and λ_φ therefore quantify how variations of the angular states influence the satisfaction of the

radial terminal condition and, indirectly, the fuel consumption through the optimality conditions.

In the unperturbed equatorial two-body configuration considered here, the mechanism that produces the required radial increase is essentially energetic and is driven by the tangential velocity component. Small perturbations of ϑ do not alter the orbital energy and do not enter the radial terminal constraint at first order. Similarly, in the equatorial case the trajectory remains confined to a fixed plane, and variations of φ do not contribute to the energy-raising mechanism responsible for achieving r_f .

Although the Hamiltonian formally contains angular terms through the spherical ZEN kinematics, the converged optimal solution remains planar and does not exploit the angular directions to satisfy the terminal condition. As a consequence, variations of ϑ and φ do not contribute at first order to the optimal transfer mechanism in this symmetric configuration, and the associated sensitivities vanish, leading to $\lambda_\vartheta(t_0) = 0$ and $\lambda_\varphi(t_0) = 0$.

A further consistency check can be obtained by relating the converged burn duration to the available acceleration level. For the baseline chemical configuration, the maximum acceleration is:

$$a_{\max} = \frac{T_{\max}}{m_0} = \frac{0.5}{462} \approx 1.08 \times 10^{-3} \text{ m/s}^2.$$

With a burn duration $\Delta t_{\text{burn}} = t_1 - t_0 \approx 25.0198$ s, a first-order estimate of the injected velocity increment is:

$$\Delta v \approx a_{\max} \Delta t_{\text{burn}} \approx 2.7 \times 10^{-2} \text{ m/s}.$$

The propellant consumption can then be recomputed using Tsiolkovsky's equation:

$$\Delta m = m_0 (1 - e^{-\Delta v/c}).$$

By substituting $c = 3000$ m/s and $\Delta v \approx 2.7 \times 10^{-2}$ m/s, one obtains:

$$\Delta m_{\text{Tsiol}} \approx 4.2 \text{ g},$$

which is in excellent agreement with the numerically obtained value of 4.170 g reported in Table 5.1. The burn duration, available thrust-to-mass ratio, and propellant expenditure are therefore mutually consistent.

It is worth noting that the ratio $\Delta m/m_0$ is extremely small in this baseline case:

$$\frac{\Delta m}{m_0} \approx \frac{0.00417}{462} \approx 9.0 \times 10^{-6}.$$

In this regime, the exponential term in Tsiolkovsky's equation can be linearized:

$$1 - e^{-\Delta v/c} \approx \frac{\Delta v}{c},$$

so that the rocket equation reduces to:

$$\Delta m \approx m_0 \frac{\Delta v}{c}.$$

This explains why the simple acceleration-based estimate of Δv and the mass consumption predicted by Tsiolkovsky's relation coincide to high accuracy in the present case.

Finally, the converged final time $t_f \approx 2913$ s indicates that the terminal constraint is satisfied after a long coast phase. Once the spacecraft has been injected onto a slightly different Keplerian trajectory through the short tangential burn, the radial coordinate reaches the prescribed value only at a specific orbital phase. The free-final-time formulation naturally captures this behaviour, leading to a solution composed of a brief energy-injection phase followed by a long drift that exploits the natural orbital dynamics.

Parameter	Value
Initial orbit altitude	600 km
Initial mass m_0	462.0 kg
Terminal radius constraint	$r_f = r_0 + 100$ m
Final time condition	$\mathcal{H}(t_f) = 0$
Switching structure	thrust-coast
Switching time t_1	25.01983188 s
Final time t_f	2913.07158287 s
Burn duration Δt_{burn}	25.01983188 s
Propellant consumed Δm	4.170 g
$\lambda_r(t_0)$	1.25099×10^{-4}
$\lambda_\vartheta(t_0)$	0
$\lambda_\varphi(t_0)$	0
$\lambda_u(t_0)$	-5.21659×10^{-4}
$\lambda_v(t_0)$	1.539973×10^{-1}
$\lambda_w(t_0)$	0
$\lambda_m(t_0)$	$9.99990974 \times 10^{-1}$

Table 5.1: Circular equatorial case – free-final-time baseline optimal solution.

5.4 Continuation to the Inclined-Orbit Case

Starting from the converged circular equatorial solution described in Section 5.3, a continuation step is performed by introducing a nonzero orbital inclination while preserving the overall formulation of the OCP. The cost functional, boundary conditions, control bounds, and indirect solution strategy remain unchanged; only the geometric configuration of the reference orbit is modified.

The rationale behind this step is to progressively increase the dynamical complexity of the scenario while maintaining numerical robustness. In the equatorial case, the motion remains confined to a symmetry plane of the two-body problem. Introducing inclination removes this planar symmetry and activates the full three-dimensional structure of the ZEN formulation, in which all state components may influence the trajectory evolution through the three-dimensional structure of the dynamics. As a consequence, the geometric symmetries that simplified the equatorial configuration no longer apply, and the optimal solution must account for fully spatial motion.

From a numerical standpoint, continuation exploits the expected smooth dependence of the optimal solution on problem parameters. Since orbital inclination enters the

dynamics continuously, the associated optimal state, costate, and switching structure are expected to vary smoothly for sufficiently small inclination increments. By using the previously converged adjoint vector as an initial guess, the indirect solver operates in a neighborhood of the new solution when solving the associated MPBVP, rather than exploring the entire admissible costate space. This substantially improves convergence reliability and reduces the risk of divergence when transitioning from a symmetric equatorial configuration to a fully three-dimensional trajectory.

The inclined-orbit case therefore represents a natural intermediate step between the planar baseline solution and more demanding scenarios, such as perturbed dynamics or combined geometric and environmental effects.

5.4.1 Free Final Time Formulation

The initial orbit remains circular, but the inclination is increased to 30° . In spherical ZEN coordinates, the initial state becomes:

$$x_0 = [r_0, \vartheta_0, \varphi_0, u_0, v_0, w_0, m_0]^\top = [R_E + 600 \text{ km}, 0, 30^\circ, 0, \sqrt{\mu/r_0} \text{ m/s}, 0, 462.0 \text{ kg}]^\top.$$

The equivalent Keplerian elements are:

$$[a, e, i, \Omega, \omega, \nu] = [r_0, 0, 30^\circ, 270^\circ, 0, 0].$$

The terminal requirement remains unchanged:

$$r_f = r_0 + 100 \text{ m}.$$

Thus, the OCP preserves exactly the same structural features as in the equatorial baseline case:

- free final time;
- terminal constraint imposed only on r_f ;
- Mayer fuel-optimal cost functional;
- thrust-coast switching structure;
- switching condition $\mathcal{S}_{\mathcal{F}}(t_1) = 0$;
- transversality condition $\mathcal{H}(t_f) = 0$;
- continuity conditions on the state at the switching times.

The only modification concerns the geometric configuration of the reference orbit. Introducing a nonzero inclination removes the planar symmetry of the equatorial case and activates the full three-dimensional structure of the ZEN formulation. All state components may now influence the trajectory evolution, although the terminal requirement remains purely radial.

The continuation strategy consists of using the converged costates of the equatorial solution as an initial guess for the inclined problem. Since inclination enters the two-body

dynamics continuously, the associated optimal solution is expected to vary smoothly for moderate inclination increments. This allows the indirect solver, applied to the associated MPBVP, to converge without altering the switching structure or reformulating the control law.

5.4.2 Converged Solution

The results for the inclined configuration are summarized in Table 5.2.

Parameter	Value
Initial orbit altitude	600 km
Inclination	30°
Initial mass m_0	462.0 kg
Terminal radius constraint	$r_f = r_0 + 100$ m
Final time condition	$\mathcal{H}(t_f) = 0$
Switching structure	thrust-coast
Switching time t_1	25.019831676 s
Final time t_f	2913.071590211 s
Burn duration Δt_{burn}	25.019831676 s
Propellant consumed Δm	4.170 g
$\lambda_r(t_0)$	1.25099×10^{-4}
$\lambda_\vartheta(t_0)$	0
$\lambda_\varphi(t_0)$	9.37×10^{-7}
$\lambda_u(t_0)$	-5.21659×10^{-4}
$\lambda_v(t_0)$	$1.53997713 \times 10^{-1}$
$\lambda_w(t_0)$	0
$\lambda_m(t_0)$	$9.99990974 \times 10^{-1}$

Table 5.2: *Inclined circular case ($i = 30^\circ$) – free-final-time baseline optimal solution.*

Several observations can be drawn from the converged solution.

First, the switching time, final time, burn duration, and propellant consumption are practically identical to those obtained in the equatorial configuration. This confirms that the maneuver remains essentially an energy-injection problem: the required radial increase is achieved through a short predominantly tangential burn, followed by a long natural Keplerian drift. The introduction of inclination does not modify the energetic nature of the terminal constraint.

Regarding the adjoint variables, λ_r , λ_u , λ_v , and λ_m remain virtually unchanged with respect to the baseline case. In particular, the dominance of $\lambda_v(t_0)$ is preserved, confirming that variations of the tangential velocity component continue to be the primary mechanism through which the terminal radial condition is satisfied at minimum propellant cost. Similarly, λ_u retains the same magnitude and sign, reflecting the secondary role of radial velocity in shaping the orbital energy variation. The mass costate λ_m remains close to unity, as expected in a Mayer fuel-optimal formulation.

The most relevant structural change concerns $\lambda_\varphi(t_0)$, which is no longer exactly zero. In the equatorial case, planar symmetry implied that variations in the φ coordinate did not influence the satisfaction of the purely radial terminal constraint at first order. Once inclination is introduced, this symmetry is removed. Although the terminal

condition still involves only r_f , the three-dimensional structure of the ZEN dynamics introduces couplings between the angular variables and the velocity components. As a consequence, small variations in φ can affect the trajectory evolution and indirectly influence the satisfaction of the endpoint condition. Consequently, the corresponding costate becomes nonzero, indicating a non-vanishing sensitivity of the optimal cost with respect to variations of the out-of-plane angular coordinate. Its magnitude, however, remains several orders of magnitude smaller than that of λ_v , confirming that the dominant mechanism of the maneuver is still energetic rather than geometric.

Finally, the persistence of the thrust-coast structure demonstrates that the optimal switching pattern is locally robust with respect to moderate inclination changes. The inclined solution is therefore smoothly connected to the equatorial baseline within the optimal-control parameter space, validating the continuation strategy adopted in the unperturbed model.

5.5 Imposition of the Selected Planar Displacement

Once the optimal solution for the inclined circular orbit has been obtained, the terminal condition is updated in order to reflect the displacement required by the collision-avoidance phase. Specifically, the purely radial terminal constraint previously imposed is replaced by a constraint on the planar displacement selected from the probabilistic collision analysis.

The modification concerns only the terminal boundary condition, while the cost functional, control bounds, dynamical model, and indirect formulation remain unchanged. The problem therefore preserves its Mayer fuel-optimal structure, with free final time and bang–bang control governed by the switching function $\mathcal{S}_{\mathcal{F}}$.

From a mathematical standpoint, the introduction of the planar displacement modifies the geometry of the terminal constraint manifold in the state space. Instead of prescribing a radial increment $r_f = r_0 + 100\text{m}$, the endpoint condition now constrains a specific in-plane displacement component, as determined by the selected avoidance direction. This alters the sensitivity of the terminal constraint with respect to the state variables and consequently modifies the associated adjoint structure through the updated terminal boundary conditions.

The continuation strategy is again employed: the converged inclined-orbit solution is used as initial guess for the updated boundary-value problem. Since the imposed displacement is moderate and enters the formulation continuously, the optimal solution is expected to continue smoothly from the previous one. Using the previously converged costates ensures that the indirect solver starts in the neighborhood of the new solution branch, thereby maintaining numerical robustness.

In this way, the progression from radial energy injection to geometrically oriented planar displacement is achieved without altering the underlying dynamical framework, allowing the impact of the modified terminal condition on the optimal control structure to be isolated and analyzed.

5.5.1 Free Final Time Formulation

The initial state is identical to that of the inclined configuration previously analyzed:

$$x_0 = \left[R_E + 600 \text{ km}, 0, 30^\circ, 0, \sqrt{\mu/r_0} \text{ m/s}, 0, 462.0 \text{ kg} \right]^T.$$

The terminal condition is now updated to:

$$r_f = r_0 + 105 \text{ m}, \quad (5.3)$$

in agreement with the displacement selected in the collision-avoidance analysis in Subsection 3.5.4.

No additional endpoint constraints are introduced. Apart from the updated radial target (5.3), the formulation of the OCP remains identical to that described in the previous sections. In particular, the free-final-time Mayer structure, the bang–bang control law, and the associated optimality conditions are unchanged.

With respect to the 100 m case, the only difference concerns the magnitude of the required terminal energy increase associated with the larger prescribed radial displacement.

5.5.2 Converged Solution

The resulting maneuver duration is:

$$t_{\text{man}} = 2912.85842378563 \text{ s}.$$

The maneuver duration t_{man} plays a central role in the subsequent development of the collision-avoidance formulation. In the present free-final-time problem, t_f emerges as part of the optimal solution and identifies the natural epoch at which the prescribed radial displacement is achieved under fuel-optimal conditions. This optimal duration therefore represents the natural time horizon of the fuel-optimal trajectory compatible with the imposed energy increase and thrust level, once the coast phase required by the orbital dynamics is taken into account.

In the following section, where the final time will be fixed and aligned with the collision epoch, the value of t_{man} provides a physically consistent reference for defining the admissible maneuver window. In particular, it indicates the characteristic time scale over which the displacement can be generated without violating optimality conditions or artificially compressing the thrust arc beyond what is dynamically feasible. The free-final-time solution thus supplies not only the optimal control profile, but also the natural temporal structure of the maneuver, which is essential for the transition to the fixed-final-time formulation.

The main characteristics of the converged solution are summarized in Table 5.3.

Several important observations can be drawn.

First, both the switching time and the propellant consumption increase slightly with respect to the 100 m inclined case.

The burn duration grows from approximately 25.02 s to 26.27 s, while the consumed propellant increases from 4.170 g to 4.378 g. This behavior is fully consistent with the larger prescribed radial displacement.

In the unperturbed two-body model, a radial increase for a near-circular orbit corresponds to an increase in specific orbital energy and therefore in semi-major axis. Since the thrust magnitude T_{\max} and exhaust velocity c are fixed, a larger required

Parameter	Value
Initial orbit altitude	600 km
Inclination	30°
Initial mass m_0	462.0 kg
Terminal displacement	105 m
Final time condition	$\mathcal{H}(t_f) = 0$
Switching structure	thrust-coast
Switching time t_1	26.270846427314 s
Final time t_f	2912.85842378563 s
Burn duration Δt_{burn}	26.270846427314 s
Propellant consumed Δm	4.378 g
$\lambda_r(t_0)$	1.25101×10^{-4}
$\lambda_\vartheta(t_0)$	0
$\lambda_\varphi(t_0)$	5.3995×10^{-8}
$\lambda_u(t_0)$	-5.12711×10^{-4}
$\lambda_v(t_0)$	$1.53998499957 \times 10^{-1}$
$\lambda_w(t_0)$	0
$\lambda_m(t_0)$	$9.99990523 \times 10^{-1}$

Table 5.3: *Inclined circular case ($i = 30^\circ$) – free-final-time optimal solution for a 105 m radial displacement.*

energy increment directly translates into a larger Δv , which in turn produces a proportional increase in burn duration and propellant mass expenditure. The consistent increase of both Δt_{burn} and Δm confirms the internal physical consistency of the solution.

Regarding the adjoint variables, the costate structure remains extremely close to that of the 100 m case. The dominant role of $\lambda_v(t_0)$ is preserved, confirming that the tangential velocity component remains the primary channel through which the terminal radial condition is satisfied at minimum propellant cost. The small variation in λ_u reflects the slightly modified balance between radial and tangential velocity contributions required to achieve the larger energy increase.

The costate $\lambda_\varphi(t_0)$ remains nonzero but extremely small. Compared to the 100 m case, its magnitude decreases slightly, indicating that the additional displacement is still achieved predominantly through energetic modification rather than through geometric reorientation of the trajectory. The remaining nonzero value reflects weak dynamical couplings introduced by the three-dimensional ZEN formulation. This confirms that, even in the inclined configuration, the optimal maneuver is essentially governed by energy variation rather than by out-of-plane geometric effects.

The vanishing of $\lambda_s(t_0)$ and $\lambda_w(t_0)$ is consistent with the previous cases and indicates that, in this configuration, the optimal cost is locally insensitive to variations of the corresponding initial state components, reflecting the residual geometric symmetries of the problem.

Finally, the very small change in the final time t_f indicates that the coast phase duration adjusts naturally to accommodate the slightly modified orbital energy, ensuring that the radial condition is satisfied at the appropriate point along the natural Keplerian evolution. The smooth variation of both the state and costate variables confirms the regularity of the solution branch and further validates the continuation strategy adopted in the unperturbed model.

5.6 Alignment with the Collision Time

The solutions obtained in the previous sections were formulated with free final time, allowing the OCP to determine autonomously the epoch at which the prescribed radial displacement is reached. This formulation is particularly useful to identify the natural maneuver duration associated with the required energy variation under fuel-optimal conditions. However, in a collision-avoidance context, the displacement must be realized at a specific physical instant, namely the effective time of closest approach identified in the probabilistic analysis.

This requirement introduces a fundamental modification in the problem structure: the terminal epoch can no longer be determined by optimality alone, but must coincide with the collision time. The maneuver must therefore be temporally aligned with an externally imposed event, while still satisfying the radial displacement constraint and preserving fuel optimality.

The transition from a free-final-time formulation to a fixed-terminal-time problem removes one degree of freedom from the boundary-value problem, corresponding to the previously free terminal epoch. In particular, the maneuver can no longer rely on a freely adjustable terminal epoch to satisfy the radial constraint; instead, the timing of the thrust arc must be selected so that the prescribed displacement is obtained exactly at the collision epoch. This leads naturally to the introduction of the maneuver start time as an additional optimization variable, while the final time is fixed.

The following subsections detail this reformulation and its implications on the structure of the optimal solution.

5.6.1 Fixed Final Time Reformulation and Optimization of the Maneuver Start

In order to guarantee that the prescribed 105 m displacement is achieved exactly at the effective collision epoch determined in the collision-avoidance analysis, the OCP is reformulated with a *fixed final time*.

Let t_{coll} denote the collision time measured with respect to the original reference

initial state. From the probabilistic collision-avoidance phase, the collision epoch is found to occur at, as explained in Subsection 3.2.2:

$$t_{\text{coll}} = 1450.273019459666 \text{ s.}$$

In the previous free-final-time formulation, the maneuver duration obtained for the 105 m displacement was:

$$t_{\text{man}} = 2912.85842378563 \text{ s.}$$

In the free-time problem, the final time emerged naturally from the optimality conditions and identified the epoch at which the prescribed radial displacement was reached under fuel-optimal conditions. In the present collision-aligned scenario, however, the terminal epoch is externally imposed and must coincide with t_{coll} . The problem is therefore reformulated by fixing the terminal time to the collision epoch.

Fixing the final time removes one degree of freedom from the boundary-value problem. Consequently, the transversality condition $\mathcal{H}(t_f) = 0$, which was previously required for free final time, no longer applies. The radial constraint:

$$r(t_f) = r_0 + 105 \text{ m}$$

remains unchanged and must now be satisfied at the prescribed collision time.

Since the displacement must still be achieved with the same thrust magnitude and exhaust velocity, the maneuver cannot adjust its terminal epoch anymore. Instead, it must be positioned in time so that the required displacement is reached exactly at t_{coll} . For this reason, the maneuver start time t_{start} is introduced as an additional optimization variable.

Because t_{start} is treated as a free temporal boundary of the OCP, the corresponding transversality condition must be satisfied at the beginning of the thrust arc. In the present Mayer formulation this yields:

$$\mathcal{H}(t_{\text{start}}) = 0.$$

This condition enforces optimality with respect to the maneuver start time and determines the correct temporal placement of the thrust arc within the trajectory.

To ensure consistency between the maneuver duration and the fixed collision epoch, an initial guess for the maneuver start time is constructed as:

$$t_{\text{start}}^{(0)} = t_{\text{coll}} - t_{\text{man}}.$$

Rather than keeping the original initial state and introducing a large temporal offset inside the boundary-value problem, the reference orbit is explicitly propagated backward in time from the original initial state to $t_{\text{start}}^{(0)}$, using the unperturbed two-body dynamics.

This choice serves two purposes. First, it provides a dynamically consistent state at the new initial epoch, avoiding the introduction of artificial discontinuities or temporal

shifts inside the shooting formulation. Second, it improves numerical robustness: the indirect method then operates over a time interval of length t_{man} , identical to that of the free-time solution, instead of integrating over a much longer arc before the maneuver begins.

The backward propagation therefore preserves the physical trajectory while redefining the temporal origin of the OCP. The procedure provides the following new initial state for the fixed-final-time formulation:

$$\begin{aligned} r_0 &= 6978000 \text{ m}, \\ \vartheta_0 &= -1.582352821489075 \text{ rad}, \\ \varphi_0 &= -0.006671897809999 \text{ rad}, \\ u_0 &= 0.0, \\ v_0 &= 6545.513200529415371 \text{ m/s}, \\ w_0 &= 3778.717355056489396 \text{ m/s}. \end{aligned}$$

After this temporal shift, the duration of the maneuver remains equal to the value previously obtained in the free-time formulation. Consequently, when time is measured from t_{start} , the fixed terminal time of the OCP becomes numerically:

$$t_f = 2912.85842378563 \text{ s},$$

ensuring that the radial displacement condition is satisfied exactly at the collision epoch.

With the final time fixed and the new initial state defined, the OCP is solved again using the indirect shooting method. The optimization variables now include:

- the initial costates $\lambda(t_0)$,
- the maneuver start time t_{start} ,
- the switching time t_1 .

The switching function condition remains enforced at the switching instant:

$$\mathcal{S}_{\mathcal{F}}(t_1) = 0,$$

while the terminal constraint imposes the prescribed radial displacement at the fixed collision epoch.

Convergence is achieved by continuation starting from the previous free-time solution, using the converged state and costate variables as an initial guess.

5.6.2 Converged Fixed-Time Solution

The main characteristics of the converged solution are summarized in Table 5.4.

Compared to the free-final-time formulation, the burn duration and propellant consumption remain essentially unchanged. The active thrust arc retains nearly the same duration as in the free-final-time case, while its temporal placement is slightly adjusted

through the optimized start time t_{start} . The propellant mass remains equal to 4.378 g within numerical precision. This confirms that the energetic requirement is dictated primarily by the prescribed 105 m radial displacement and not by the temporal placement of the maneuver. Fixing the terminal epoch does not modify the amount of orbital energy that must be injected into the system; it only constrains the orbital phase at which this energy increment must be realized.

Parameter	Value
Initial orbit altitude	600 km
Inclination	30°
Initial mass m_0	462.0 kg
Terminal displacement	105 m
Final time condition	t_f fixed, $\mathcal{H}(t_{\text{start}}) = 0$
Switching structure	thrust-coast
Start time t_{start}	-0.022967726492 s
Switching time t_1	26.247879085273 s
Final time t_f	2912.85842378563 s
Burn duration Δt_{burn}	26.270846811765 s
Propellant consumed Δm	4.378 g
$\lambda_r(t_0)$	1.25101×10^{-4}
$\lambda_s(t_0)$	0
$\lambda_\varphi(t_0)$	9.676229×10^{-8}
$\lambda_u(t_0)$	$-5.13670657 \times 10^{-4}$
$\lambda_v(t_0)$	$1.33370548896 \times 10^{-1}$
$\lambda_w(t_0)$	$7.6994545902 \times 10^{-2}$
$\lambda_m(t_0)$	$9.99990523010 \times 10^{-1}$

Table 5.4: *Inclined circular case ($i = 30^\circ$) – fixed-final-time optimal solution aligned with the collision epoch.*

The optimized start time t_{start} is very close to zero, with a small negative value on the order of 10^{-2} seconds. This indicates that the collision epoch and the free-time optimal terminal epoch were already nearly aligned. The backward-propagation procedure therefore introduces only a minimal temporal shift, and the solution remains on the same regular branch previously identified.

The most relevant differences with respect to the free-final-time case appear in the velocity costates. While $\lambda_r(t_0)$ and $\lambda_m(t_0)$ remain practically unchanged, the relative magnitudes of $\lambda_v(t_0)$ and $\lambda_w(t_0)$ are modified: λ_v decreases and λ_w becomes clearly non-zero.

In the adopted ZEN frame, u denotes the radial component, v the eastward component, and w the northward component. For an inclined orbit, the spacecraft velocity generally possesses both eastward and northward components, whose relative magnitudes depend on the geocentric longitude and latitude of the spacecraft. After the backward temporal shift associated with t_{start} , the maneuver begins at a different point along the nominal orbit. At this new initial epoch, the velocity is no longer predominantly aligned with the eastward direction; instead, it exhibits a visible northward component w_0 .

Since the velocity costates determine the direction of the optimal thrust through the primer vector, the components λ_v and λ_w must adjust so as to produce the thrust orientation that is optimal at the beginning of the burn. In the present case, this optimal

direction remains nearly tangent to the local orbit, and therefore reflects the local decomposition of the spacecraft velocity into eastward and northward components. This was reflected in the dominance of λ_v and the negligible value of λ_w .

In the present fixed-time solution, the maneuver starts at a slightly different orbital phase, where both eastward and northward velocity components are present. To generate a thrust direction that remains effectively tangent to the orbit at this new location, both λ_v and λ_w must be non-zero and assume appropriate relative magnitudes. The observed redistribution between these two costates therefore reflects the geometric change in the initial velocity direction.

Importantly, this adjustment does not modify the energetic nature of the maneuver. The burn duration and propellant mass remain practically identical, confirming that the required orbital energy increment is unchanged. What changes is the decomposition of the optimal thrust direction within the local East–North subspace, consistent with the new starting point along the inclined orbit.

The smooth variation of the costates and of the maneuver parameters confirms that the fixed-final-time solution remains continuously connected to the free-time branch. The reformulation affects the temporal placement of the maneuver but preserves its fundamental structure and energetic mechanism.

5.6.3 Angular Consistency Verification

As a final verification step, the angular components of the terminal state are compared with those of the nominal reference trajectory at the collision epoch.

In the collision-avoidance analysis, the imposed displacement was purely radial at the collision time. This means that, at t_{coll} , the optimized trajectory was required to differ from the nominal one only in the radial coordinate r , while the angular coordinates ϑ (longitude) and φ (geocentric latitude) were expected to remain unchanged. In other words, the maneuver was designed to generate a separation along the local Zenith direction, without introducing any deliberate shift along the East or North directions of the ZEN frame.

Because the terminal time is fixed and the maneuver start time is optimized, it is essential to verify that the solution does not introduce unintended angular deviations as a side effect of the thrusting phase. Any significant discrepancy in ϑ or φ at $t_f = t_{\text{coll}}$ would indicate that part of the control action had been used to alter the angular position, thereby violating the purely radial nature of the required displacement.

The residual differences between the optimized and nominal angular coordinates at the fixed terminal time are found to be:

$$\Delta\vartheta \approx 10^{-5} \text{ rad}, \quad \Delta\varphi \approx 10^{-5} \text{ rad}.$$

These correspond approximately to:

$$10^{-5} \text{ rad} \approx 5.7 \times 10^{-4} \text{ deg},$$

which is well below one millidegree.

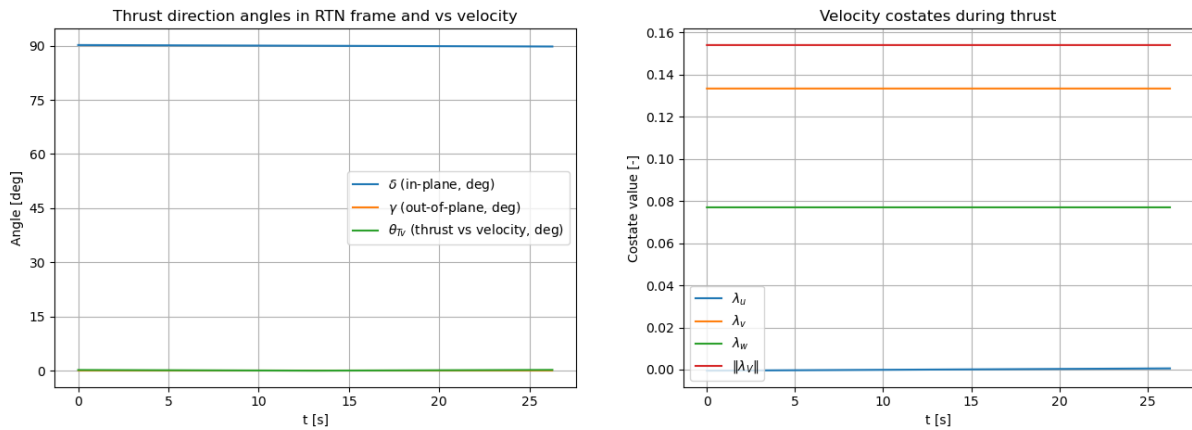
The extremely small magnitude of these residuals confirms the angular consistency of the optimized trajectory. The maneuver therefore achieves the required collision-avoidance displacement exclusively along the radial direction at the collision epoch, without introducing meaningful longitudinal or latitudinal shifts. This result validates both the fixed-final-time reformulation and the optimization of the maneuver start time, demonstrating that the imposed boundary conditions are satisfied with high precision and that the geometric interpretation of the maneuver remains fully coherent.

5.6.4 Thrust Direction and Velocity Costates Behavior

To further characterize the optimal solution, the thrust direction during the active burn arc and the corresponding evolution of the velocity costates are analyzed.

Figure 5.1 reports:

- the thrust direction angles expressed in the RTN reference frame;
- the angle between the thrust vector and the instantaneous velocity vector;
- the time evolution of the velocity costates $\lambda_u, \lambda_v, \lambda_w$ during the thrust phase.



5.1a Thrust direction angles in the RTN frame and thrust–velocity angle during the burn arc.

5.1b Velocity costates $\lambda_u, \lambda_v, \lambda_w$ and their norm during the burn arc.

Figure 5.1: Thrust direction and velocity costate behavior during the active thrust phase.

As discussed in Section 4.2.3, the optimal thrust direction is aligned with the normalized velocity costate vector (primer vector):

$$\hat{\mathbf{t}}^* = \frac{\boldsymbol{\lambda}_V}{|\boldsymbol{\lambda}_V|}, \quad \boldsymbol{\lambda}_V = [\lambda_u, \lambda_v, \lambda_w]^\top.$$

Therefore, the thrust pointing law is entirely governed by the *relative* magnitudes of the velocity costates: only the ratios among $(\lambda_u, \lambda_v, \lambda_w)$ determine the direction of $\hat{\mathbf{t}}^*$, while any common scaling of $\boldsymbol{\lambda}_V$ leaves the pointing unchanged.

From the right-hand plot of Figure 5.1, it can be observed that during the burn arc the components λ_u , λ_v , and λ_w remain almost constant. This behavior is consistent with the relatively short burn duration (approximately 26 s) compared to the orbital period, which is of the order of several thousand seconds. Over such a limited interval, the adjoint dynamics evolves only slightly, and the velocity costate vector undergoes negligible rotation in its three-dimensional space.

As a direct consequence, the direction of the primer vector remains nearly constant throughout the burn, which explains the almost constant thrust angles observed in the left-hand plot. In particular:

- the in-plane thrust angle remains essentially constant. This angle is measured in the orbital plane starting from the radial direction \hat{R} toward the tangential direction \hat{T} . In the present case it is approximately 90° , indicating that the thrust is applied almost entirely along the tangential direction;
- the out-of-plane component is negligible, confirming that the thrust has no normal component and that the maneuver preserves the orbital plane throughout the burn;
- the angle between thrust and velocity is nearly constant and very small, showing that the thrust direction is almost aligned with the instantaneous velocity vector.

The predominance of the tangential direction is fully consistent with the physical objective of the maneuver. A small increase in orbital radius for a near-circular orbit is achieved most efficiently by increasing the specific orbital energy. For quasi-circular motion, a variation of the semi-major axis is predominantly driven by changes in the tangential component of the velocity. Applying thrust nearly parallel to the instantaneous velocity therefore produces the required increase in orbital energy while avoiding unnecessary radial components that would not contribute efficiently to the desired semi-major axis variation. In this sense, the predominance of the tangential thrust direction reflects the energetic optimality of the maneuver.

The velocity costates provide the corresponding optimal-control interpretation of this mechanism. Since \hat{t}^* is collinear with λ_V , the components $(\lambda_u, \lambda_v, \lambda_w)$ define how the optimal policy distributes thrust between the local radial direction and the two local transverse directions of the ZEN frame (eastward v and northward w). In particular, the small magnitude of λ_u indicates that the optimal burn avoids introducing a significant radial thrust component, consistently with the energetic interpretation above. At the same time, the relative balance between λ_v and λ_w must be interpreted in a *geometric* sense: for an inclined orbit, the instantaneous velocity generally has both eastward and northward components, and their relative weights depend on the orbital location at which the maneuver is initiated (geocentric longitude/latitude and the corresponding local orientation of the velocity in the ZEN basis).

Accordingly, the distribution between λ_v and λ_w reflects the local decomposition of the direction that is optimal for energy injection within the moving ZEN frame at the chosen start epoch. When the maneuver starts from a point where the velocity is mostly eastward, the primer vector tends to be dominated by the v -component and λ_v becomes larger. Conversely, when the maneuver is initiated from a point where a significant

portion of the velocity lies along the northward direction, the primer vector must incorporate a corresponding w -component, and λ_w becomes comparably relevant. In all cases, what matters is that the resulting $\hat{\mathbf{t}}^*$ remains close to the instantaneous velocity direction, consistently with a fuel-optimal energy-injection strategy for a small-radius increase.

Finally, the near-constancy of the velocity costates and of the thrust direction during the active arc confirms the internal consistency of the indirect solution. The bang–bang structure is preserved, the switching occurs cleanly when the switching function vanishes, and no evidence of additional switches or singular behavior is observed over the burn interval. Altogether, these results validate both the optimality conditions derived from Pontryagin’s Maximum Principle and their numerical enforcement within the shooting-based implementation.

5.7 Electric-Propulsion Cases

After obtaining a first converged fixed-final-time solution with chemical propulsion, characterized by a short burn duration and comparatively larger propellant expenditure, the same unperturbed optimal-control problem is solved again by modifying only the propulsion parameters. This step constitutes a parametric continuation in (T_{\max}, c) , aimed at assessing how the optimal control law adapts when the thrust level and exhaust velocity are representative of electric propulsion systems.

The dynamical model, boundary conditions, terminal constraints, and switching structure remain unchanged. In particular, the prescribed radial displacement at the fixed collision epoch is kept identical to the chemical baseline case. Only the maximum thrust T_{\max} and the exhaust velocity c are varied, thereby isolating the effect of propulsion technology on the optimal solution.

Two electric-propulsion configurations are considered:

- a configuration representative of ion propulsion, characterized by very low thrust and high exhaust velocity;
- a configuration representative of Hall-effect propulsion, characterized by higher thrust than the ion case and moderately high exhaust velocity.

From a theoretical standpoint, the transition from chemical to electric propulsion modifies the relative importance of thrust magnitude and propellant efficiency. Lower thrust levels imply a weaker instantaneous control authority, while higher exhaust velocity reduces the propellant mass required to achieve a given velocity increment. Within a fixed final-time framework, this combination is expected to produce a significantly different distribution of thrust over the maneuver interval.

Compared to the chemical baseline, the electric-propulsion cases generally exhibit:

- a longer active thrust interval, often extending over a substantial portion of the available time horizon, reflecting the reduced acceleration capability;

- a lower propellant consumption due to the higher exhaust velocity, despite the longer burn duration;
- a more pronounced time variation of the velocity costates during the burn, leading to a less constant thrust direction in accordance with the optimal pointing law and with the dependence of $\hat{\mathbf{t}}^*$ on the primer vector λ_V .

The following subsections present the detailed results for each configuration and analyze how the structure of the optimal solution evolves as the propulsion parameters move from the impulsive-leaning chemical regime toward the low-thrust electric regime.

5.7.1 Hall-Effect Thruster Case: $T_{\max} = 0.05$ N, $c = 20000$ m/s

The first electric-propulsion configuration considered is representative of a Hall-effect thruster, characterized by a thrust level significantly lower than the chemical baseline and by a high exhaust velocity typical of electric propulsion systems. The dynamical model, terminal constraints, and fixed-final-time condition remain unchanged; only the propulsion parameters (T_{\max}, c) are modified.

Parameter	Value
Propulsion type	Hall-effect thruster
T_{\max}	0.05 N
c	20000 m/s
Initial mass m_0	462.0 kg
Switching structure	thrust-coast
Terminal displacement	105 m
Final time t_f	2912.881391512122 s
Start time t_{start}	237.072591110485 s
Switching time t_1	509.153559990403 s
Burn duration Δt_{burn}	272.080969 s
Propellant consumed Δm	0.680 g
$\lambda_r(t_0)$	1.980584×10^{-5}
$\lambda_\vartheta(t_0)$	0
$\lambda_\varphi(t_0)$	-1.37944×10^{-7}
$\lambda_u(t_0)$	$1.48370686 \times 10^{-3}$
$\lambda_v(t_0)$	$2.1164417381 \times 10^{-2}$
$\lambda_w(t_0)$	$1.1859089292 \times 10^{-2}$
$\lambda_m(t_0)$	$9.99998484678 \times 10^{-1}$

Table 5.5: *Inclined circular case ($i = 30^\circ$) – fixed-final-time solution with Hall-effect propulsion.*

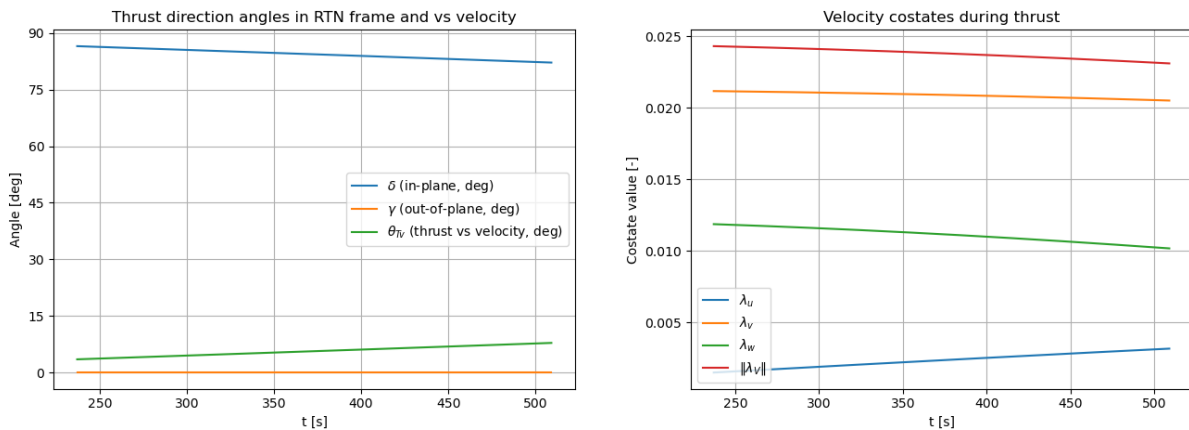
Several structural differences with respect to the chemical-propulsion baseline immediately emerge.

First, the burn duration increases by more than one order of magnitude, from approximately 26 s in the chemical case to about 272 s in the Hall-effect configuration. This extension of the active arc is a direct consequence of the lower thrust level. Since the required radial displacement at the prescribed collision epoch is unchanged, the overall variation in orbital energy — and therefore in semi-major axis — remains essentially the same as in the chemical case. However, the reduced thrust authority implies that the same impulse must be distributed over a much longer time interval.

Second, the propellant consumption drops dramatically from several grams in the chemical case to 0.680 g. This reduction is entirely consistent with the higher exhaust velocity. Although the burn duration is longer, the mass flow rate $\dot{m} = T_{\max}/c$ is substantially smaller, leading to a lower total mass expenditure for the same overall orbital-energy modification. The fixed-final-time constraint does not alter the required energy increment, but the higher c allows that increment to be achieved more efficiently in terms of propellant mass.

A third important difference concerns the optimal maneuver placement. The optimized start time shifts to $t_{\text{start}} \approx 237$ s, which is significantly different from the chemical baseline. In the low-thrust regime, the optimizer must select a maneuver window that maximizes the effectiveness of the prolonged thrust arc under the fixed terminal constraint. In the chemical solution, the short burn behaves almost impulsively and its effectiveness is weakly sensitive to the exact orbital phase at which it is applied. In contrast, in the low-thrust Hall configuration the acceleration is distributed over several minutes, so its cumulative effect depends more strongly on the orbital geometry during the burn. The optimizer therefore shifts t_{start} to a region of the orbit where the distributed tangential acceleration produces the most favorable secular variation in semi-major axis under the fixed terminal constraint.

The longer burn also affects the adjoint dynamics. Unlike the chemical case, where the short thrust arc led to nearly constant velocity costates, here the burn duration represents a non-negligible fraction of the orbital time scale. Consequently, the components of λ_V exhibit a more visible time evolution during the active phase. Since the thrust direction is aligned with the normalized velocity costate vector, this induces a measurable rotation of the thrust direction throughout the burn, as shown in Figure 5.2.



5.2a Thrust-direction angles in RTN and thrust-velocity angle during the burn.

5.2b Velocity costates and their norm during the burn.

Figure 5.2: Thrust direction and velocity costate behavior during the active thrust phase. Electric-propulsion case with $T_{\max} = 0.05$ N and $c = 20000$ m/s.

In contrast with the nearly time-invariant thrust direction observed in the chemical solution, the Hall-effect case displays a progressive and continuous reorientation of the thrust vector during the burn. This behavior is a direct consequence of the longer active-

control interval. Since the thrust remains applied for several minutes, the direction determined by the velocity costates influences the state over a significantly extended time span. Even moderate variations of λ_V during the burn therefore translate into a gradual reorientation of the thrust vector. In contrast, in the chemical case the very short burn interval limited the effect of costate evolution, resulting in an almost constant thrust direction.

The bang–bang structure is nevertheless preserved, and the switching remains unique and well defined. No additional switching points or singular arcs appear in the low-thrust regime, confirming that the qualitative structure of the optimal policy is unchanged, while only its temporal distribution and geometric realization are modified by the propulsion parameters.

5.7.2 Ion Thruster Case: $T_{\max} = 0.01$ N, $c = 30000$ m/s

The second electric-propulsion configuration is representative of ion propulsion, characterized by a further reduction of the maximum thrust and by an even higher exhaust velocity. As in the Hall-effect case, the dynamical model, boundary conditions, and fixed-final-time constraint remain unchanged; only the propulsion parameters (T_{\max}, c) are modified.

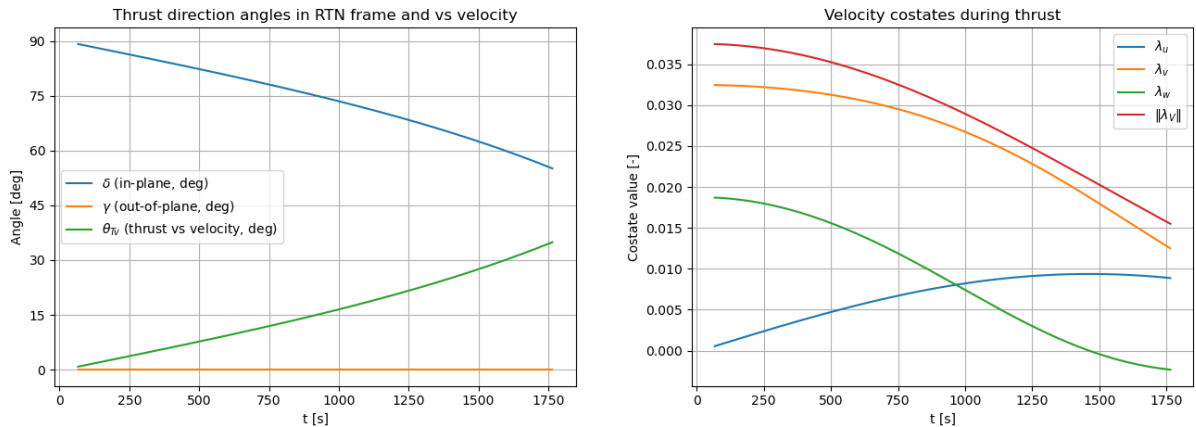
Parameter	Value
Propulsion type	Ion thruster
T_{\max}	0.01 N
c	30000 m/s
Initial mass m_0	462.0 kg
Switching structure	thrust-coast
Terminal displacement	105 m
Final time t_f	2912.881391512122 s
Start time t_{start}	66.836423265969 s
Switching time t_1	1770.642704052697 s
Burn duration Δt_{burn}	1698.108049 s
Propellant consumed Δm	0.568 g
$\lambda_r(t_0)$	3.0424839×10^{-5}
$\lambda_\vartheta(t_0)$	0
$\lambda_\varphi(t_0)$	1.2732×10^{-8}
$\lambda_u(t_0)$	$5.52681818 \times 10^{-4}$
$\lambda_v(t_0)$	$3.2440560993 \times 10^{-2}$
$\lambda_w(t_0)$	$1.8696952771 \times 10^{-2}$
$\lambda_m(t_0)$	$9.99997693759 \times 10^{-1}$

Table 5.6: *Inclined circular case ($i = 30^\circ$) – fixed-final-time solution with ion propulsion.*

The most evident difference with respect to the chemical baseline is the dramatic increase in burn duration. The active thrust arc extends over approximately 1700 s, representing a substantial fraction of the available time horizon. This behavior is a direct consequence of the extremely low thrust level. Since the prescribed radial displacement at the fixed collision epoch remains unchanged, the required variation in orbital energy is essentially the same as in the chemical case. However, the much lower thrust authority implies that the corresponding velocity increment must be accumulated gradually over a long time interval.

Despite the very long burn, the propellant consumption decreases further to 0.568 g. This reduction is entirely consistent with the high exhaust velocity. Although thrust is applied over a large portion of the time horizon, the mass flow rate $\dot{m} = T_{\max}/c$ is extremely small, and the overall propellant expenditure remains lower than in the Hall-effect configuration and dramatically lower than in the chemical baseline. The fixed-final-time formulation therefore highlights clearly the trade-off between thrust magnitude and propellant efficiency: lower thrust stretches the maneuver in time, while higher exhaust velocity reduces the mass cost of achieving the same orbital-energy modification.

The optimal start time is again shifted with respect to the chemical solution, taking the value $t_{\text{start}} \approx 67$ s. In the ion case, the thrust arc occupies a significant portion of the available horizon; its cumulative effect is therefore strongly dependent on the orbital phase over which it is distributed. The optimizer selects a maneuver window that ensures that the prolonged tangential acceleration produces the required secular increase in semi-major axis exactly at the prescribed collision epoch. Compared to the chemical case, where the short burn behaves almost impulsively, the ion solution is intrinsically phase-sensitive: the timing of the burn window plays a crucial role in satisfying the fixed terminal constraint.



5.3a Thrust-direction angles in RTN and thrust-velocity angle during the burn.

5.3b Velocity costates and their norm during the burn.

Figure 5.3: Thrust direction and velocity costate behavior during the active thrust phase. Electric-propulsion case with $T_{\max} = 0.01$ N and $c = 30000$ m/s.

The very long active-control interval also has a visible impact on the adjoint evolution. Over a burn lasting nearly half an orbital period, the velocity costates λ_u , λ_v , and λ_w evolve appreciably in time. Since the optimal thrust direction is aligned with the normalized velocity costate vector, this evolution induces a continuous and more pronounced rotation of the thrust vector during the burn, as shown in Figure 5.3. The variation of the thrust-direction angles is therefore a direct manifestation of the gradual evolution of the primer vector over the extended control interval.

Nevertheless, the qualitative structure of the solution remains unchanged. The bang-bang character of the control is preserved, the switching remains unique, and no

singular arcs appear despite the extremely low thrust level. The maneuver is strictly planar, as indicated by the negligible normal component of the thrust throughout the burn. The differences with respect to the chemical baseline therefore concern primarily the temporal distribution of the control and the progressive reorientation of the thrust direction, rather than the fundamental structure of the optimal policy.

5.7.3 Comparative Discussion

A global comparison among the chemical, Hall-effect, and ion-propulsion configurations highlights the fundamental trade-off between thrust authority and exhaust velocity within a fixed-final-time optimal-control framework.

In the chemical baseline, the relatively high thrust level allows the required orbital-energy increment to be delivered over a very short time interval. The active thrust arc behaves almost impulsively, the velocity costates remain nearly constant during the burn, and the thrust direction appears almost time-invariant. The maneuver is therefore concentrated in time, weakly sensitive to the exact orbital phase, and characterized by a larger propellant expenditure due to the comparatively low exhaust velocity.

Transitioning to electric propulsion fundamentally modifies the temporal distribution of the control. Both electric configurations exhibit substantially longer burn durations, reflecting their lower thrust authority. Since the prescribed radial displacement at the collision epoch is identical in all cases, the required variation in orbital energy remains essentially unchanged. What changes is the way in which that energy increment is delivered: gradually rather than impulsively. As the thrust arc lengthens, the evolution of the velocity costates becomes more visible over the active interval and the thrust direction, being aligned with the primer vector, undergoes a more gradual and more pronounced rotation.

Comparing the two electric cases, the ion configuration represents the extreme low-thrust regime. With $T_{\max} = 0.01$ N, the burn occupies a large fraction of the available horizon, and the maneuver becomes strongly phase-dependent. The prolonged action of the control amplifies the influence of the time evolution of λ_V on the thrust direction, resulting in a more evident reorientation of the control vector during the burn. At the same time, the very high exhaust velocity minimizes the propellant mass required to achieve the prescribed orbital-energy variation.

The Hall-effect configuration occupies an intermediate position between the chemical and ion limits. Its thrust level is sufficient to reduce the burn duration by roughly one order of magnitude compared to the ion case, while still remaining significantly longer than the chemical burn. Accordingly, the variation of the velocity costates and the rotation of the thrust direction are more pronounced than in the chemical solution but less extreme than in the ion case. In terms of propellant mass, it also provides an intermediate performance, reflecting its exhaust velocity between the chemical and ion extremes.

In all three configurations, the qualitative structure of the optimal solution remains unchanged. The control retains a bang-bang character, a single switching point is ob-

served, and no singular arcs appear even in the very low-thrust regime. The maneuver remains strictly planar in all cases, consistent with the purely radial displacement requirement of the unperturbed collision-avoidance setup.

An additional aspect that differentiates the impulsive-like chemical solution from the low-thrust regimes concerns the duration over which the control remains active. When the thrust arc extends over a large fraction of the time horizon, as in the ion configuration, the direction prescribed by the velocity costates acts on the state dynamics for a much longer interval. As a result, even moderate variations in λ_V during the burn produce a cumulative effect on the trajectory through the control law. This explains the increased sensitivity of the solution to the time evolution of the velocity costates in the ion case: the prolonged action of the control makes the gradual rotation of the primer vector dynamically significant, whereas in the short chemical burn its influence is confined to a very limited time span.

Overall, the comparison demonstrates that varying (T_{\max}, c) does not alter the fundamental optimal-control structure of the problem, but it significantly affects the temporal distribution of the control, the phase sensitivity of the maneuver window, and the time evolution of the thrust-direction angles. The chemical case approaches an impulsive limit, while the ion configuration approaches a continuous low-thrust regime, with the Hall-effect solution providing a transition between these two extremes.

6 Influence of Perturbations on the Optimal Solution

This chapter investigates how the optimal collision-avoidance maneuver is modified when environmental perturbations are progressively incorporated into the dynamical model. While the previous chapter considered the ideal two-body setting, real orbital dynamics is affected by several perturbing forces that alter both the spacecraft trajectory and the optimal-control structure of the maneuver. In order to assess these effects in a systematic way, the model fidelity is increased in two steps: first by introducing third-body gravitational perturbations due to the Sun and the Moon, and subsequently by including the effects of Earth non-sphericity.

The optimal-control formulation remains unchanged with respect to the unperturbed reference case. The maneuver must still achieve the prescribed radial displacement at the collision epoch while minimizing propellant consumption, under the same propulsion parameters and switching structure adopted previously. Maintaining the same optimal-control framework allows a direct comparison with the two-body solution, so that any modification in the optimal trajectory or in the adjoint variables can be attributed to the changes induced in the dynamical system by the additional perturbing forces.

A fundamental conceptual difference with respect to the ideal two-body problem arises from the explicit dependence of the perturbed dynamics on the *absolute epoch*. Third-body accelerations depend on the instantaneous ephemerides of the Sun and the Moon, while Earth non-sphericity depends on the spacecraft position relative to the rotating Earth-fixed gravity field. As a consequence, the dynamical system becomes explicitly time dependent, and the time-translation invariance that characterizes the ideal circular two-body model is lost.

In the unperturbed setting, selecting a different starting point along the same circular orbit leads to an equivalent dynamical configuration. The velocity magnitude is uniquely determined by the orbital radius and the equations of motion are autonomous, so a shift of the initial anomaly simply corresponds to a rigid rotation of the same orbit. In the perturbed case, instead, the instantaneous radius and velocity components vary along the trajectory due to periodic and secular effects generated by the perturbing accelerations. The dynamical evolution therefore depends not only on the geometric position along the orbit, but also on the absolute time at which the maneuver is initiated.

For this reason, both the backward propagation step and the indirect correction procedure must preserve absolute-time consistency in order to apply perturbations coherently.

The overall solution strategy therefore follows the same fixed-final-time formulation adopted in the unperturbed case:

- the collision epoch is treated as a prescribed final time t_f in the optimal-control problem;
- backward propagation is performed to construct a dynamically consistent initial state and to ensure that the maneuver horizon is sufficiently long to achieve the required terminal displacement;
- the OCP is solved with fixed t_f , optimizing the maneuver start time t_{start} and the switching time(s) associated with the selected structure.

In the following sections, this procedure is applied first to the lunisolar-only perturbation model and then to the combined lunisolar plus Earth non-sphericity case. This progressive approach allows a systematic assessment of how increasing dynamical fidelity modifies both the optimal trajectory and the associated adjoint structure, while preserving the same optimal-control formulation and terminal requirements.

6.1 Lunisolar Perturbations Case

6.1.1 Backward Propagation

When third-body gravitational perturbations are included, the spacecraft state at the beginning of the maneuver cannot be constructed by a simple time shift of the unperturbed circular reference. In this higher-fidelity setting, the equations of motion depend explicitly on the absolute epoch through the Sun–Moon ephemerides, and therefore the initial condition must be consistent with the physical time at which the maneuver is initiated. For this reason, the state at the start of the optimal-control horizon is obtained by backward propagation from the prescribed collision epoch, using the same lunisolar-perturbed dynamics adopted in the forward integration. This guarantees that the state used as initial boundary condition corresponds to the correct absolute configuration of the spacecraft, Earth, and third bodies.

The resulting initial state in ZEN coordinates is:

$$\begin{aligned}
 r_0 &= 6978000.330589165 \text{ m}, \\
 \vartheta_0 &= -1.582352641062560, \\
 \varphi_0 &= -0.006671931049932, \\
 u_0 &= 0.000309455337941, \\
 v_0 &= 6545.512097161420 \text{ m/s}, \\
 w_0 &= 3778.717575725471 \text{ m/s}.
 \end{aligned} \tag{6.1}$$

Two remarks follow from (6.1). First, the radius differs slightly from the ideal two-body value, and a small but nonzero radial velocity component is present ($u_0 \neq 0$). This is physically expected: even when the nominal orbit is close to circular, lunisolar perturbations introduce weak periodic and secular variations of the orbital elements, leading to small oscillations of the radial distance and therefore to a nonzero \dot{r} . In

the two-body circular model the equilibrium condition enforces $u \equiv 0$ at all times, whereas the additional third-body acceleration breaks this exact balance and makes the instantaneous orbit only approximately circular.

Second, the tangential velocity components v_0 (eastward) and w_0 (northward) remain of the same order as in the unperturbed reference case, confirming that the motion is still predominantly tangential and that the trajectory remains close to the original LEO configuration. The small magnitude of u_0 relative to v_0 and w_0 indicates that the orbit departs only mildly from circularity over the considered time window, but this mild departure is sufficient to break the strict orbit-equivalence property of the two-body case. As a consequence, enforcing the backward propagation with the correct absolute epoch is essential to initialize the fixed-final-time OCP consistently and to ensure that the subsequent optimal solution is computed for the actual perturbed state corresponding to the maneuver start epoch.

6.1.2 Converged Solution

Solving the fixed-final-time OCP with the same boundary conditions and switching structure yields the following converged solution:

Parameter	Value
Perturbation model	Lunisolar only
Switching structure	thrust-coast
Terminal displacement	105 m (radial)
Start time t_{start}	-0.024979896655 s
Switching time t_1	26.244909917777 s
Final time t_f	2912.85842378563 s
Burn duration Δt_{burn}	26.269889814432 s
Propellant consumed Δm	4.378 g
$\lambda_r(t_0)$	1.25101×10^{-4}
$\lambda_\vartheta(t_0)$	$-8.1999221 \times 10^{-5}$
$\lambda_\varphi(t_0)$	$-1.9122354 \times 10^{-5}$
$\lambda_u(t_0)$	$-5.13670657 \times 10^{-4}$
$\lambda_v(t_0)$	$1.33370548896 \times 10^{-1}$
$\lambda_w(t_0)$	$7.6994545902 \times 10^{-2}$
$\lambda_m(t_0)$	$9.99990523010 \times 10^{-1}$

Table 6.1: *Inclined circular case ($i = 30^\circ$) – fixed-final-time solution including lunisolar perturbations.*

A first important observation concerns the overall structure of the optimal solution. Despite the inclusion of third-body perturbations, the optimal-control policy retains the same qualitative structure obtained in the unperturbed chemical-propulsion case. The switching structure remains thrust-coast, a single switching point is observed, and the duration of the active thrust arc remains extremely close to the value obtained in the two-body reference solution. The burn duration is approximately 26.27 s, which differs only marginally from the unperturbed case. Similarly, the propellant consumption remains essentially unchanged, with a total mass expenditure of about 4.38 g. This result confirms that, over such a short maneuver interval, the integrated effect of lunisolar perturbations on the spacecraft energy balance is relatively small. Consequently, the total impulse required to achieve the prescribed radial displacement at the collision epoch remains almost identical to the two-body case.

The switching time also remains very close to the value obtained in the unperturbed solution. This indicates that the optimal placement of the thrust arc within the fixed-final-time horizon is only weakly affected by the presence of the third-body accelerations. Since the maneuver duration represents a very small fraction of the orbital period, the perturbing accelerations act essentially as a slowly varying background disturbance during the burn. As a result, the optimizer does not need to significantly modify the temporal distribution of the control in order to satisfy the terminal radial constraint.

A more visible difference emerges in the structure of the adjoint variables. In contrast with the unperturbed two-body case, the angular costates λ_ϑ and λ_φ are no longer negligible. Their presence reflects the loss of rotational symmetry introduced by the lunisolar perturbations. In the ideal two-body circular model the dynamics is invariant under rotations along the orbit, and therefore the optimal solution for a purely radial terminal constraint does not involve any significant sensitivity with respect to the angular coordinates. Once third-body perturbations are included, this symmetry is broken. The perturbing acceleration depends on the instantaneous configuration of the spacecraft relative to the Sun and the Moon, which introduces an explicit dependence of the dynamics on the orbital phase and on the absolute epoch. As a consequence, variations in the angular coordinates influence the subsequent dynamical evolution and therefore appear in the optimality conditions through nonzero angular costates.

The magnitude of these angular costates nevertheless remains significantly smaller than the dominant velocity costates. In particular, the components λ_v and λ_w still govern the thrust pointing law through the primer vector λ_V , exactly as in the unperturbed case. This indicates that the maneuver remains primarily driven by the same physical mechanism identified previously: a tangential acceleration that increases the orbital energy in order to produce the required variation of the semi-major axis. The additional angular sensitivities therefore represent a secondary correction associated with the weak coupling introduced by the third-body perturbations.

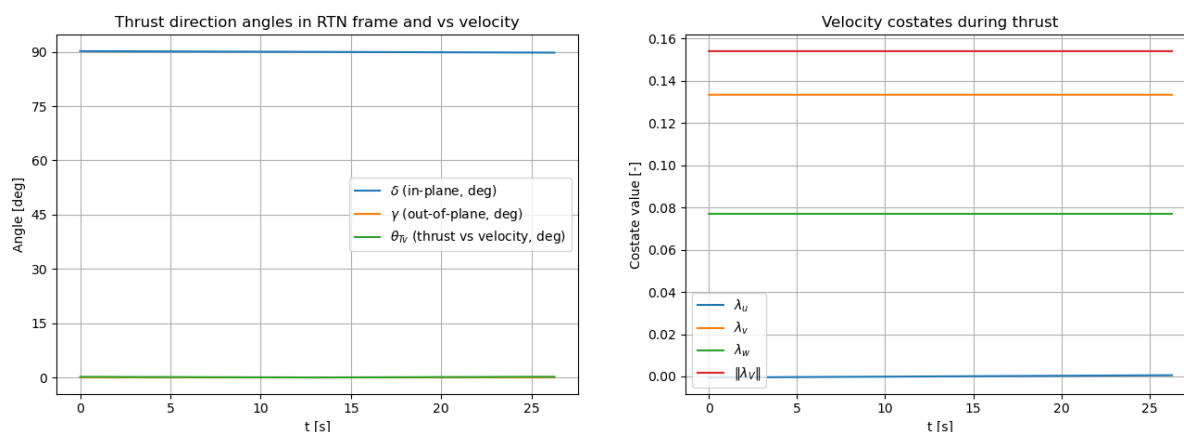
The small negative value of the optimal start time, $t_{\text{start}} \approx -0.025$ s, is also consistent with the unperturbed solution and remains essentially negligible compared with the maneuver duration. This value indicates that the optimal thrust arc begins almost immediately at the start of the control horizon, and the slight shift simply reflects the numerical adjustment required to satisfy the terminal constraint exactly in the perturbed dynamical environment.

Overall, the comparison with the unperturbed chemical-propulsion solution shows that lunisolar perturbations introduce measurable modifications in the adjoint structure of the optimal solution while leaving the main energetic characteristics of the maneuver essentially unchanged. The thrust arc duration, propellant consumption, and switching structure remain nearly identical to the two-body case, whereas the appearance of nonzero angular costates highlights the loss of dynamical symmetry and the explicit dependence of the optimal solution on the absolute epoch.

6.1.3 Thrust Direction and Velocity Costates Behavior with Lunisolar Perturbations

To further characterize the structure of the optimal solution in the presence of lunisolar perturbations, the thrust direction angles in the RTN frame and the evolution of the velocity costates during the burn arc are analyzed.

As in the unperturbed case, the thrust direction is reconstructed from the primer vector $\lambda_V = [\lambda_u, \lambda_v, \lambda_w]^T$, according to the optimal pointing law discussed in Section 4.2.3. The time evolution of λ_V directly determines the orientation of the thrust vector.



6.1a Thrust-direction angles in the RTN frame and thrust-velocity angle during the burn.

6.1b Velocity costates and their norm during the burn.

Figure 6.1: Thrust direction and velocity costate behavior during the active thrust phase. Lunisolar perturbations case.

The results exhibit a behaviour qualitatively very similar to the unperturbed short-burn solution obtained with chemical propulsion. In particular, the velocity costates remain almost constant over the burn duration (approximately 26 s). Since the thrust arc is extremely short compared both to the orbital period and to the characteristic time scales associated with the lunisolar perturbations, the perturbing accelerations introduce only very small variations in the adjoint dynamics during this interval. As a consequence, the velocity costate vector λ_V undergoes only negligible rotations over the thrust arc.

This behaviour is fully consistent with the structure observed in the unperturbed case. In the chemical-propulsion solution without perturbations, the burn duration was already very short relative to the orbital time scale, which resulted in nearly constant velocity costates and therefore in an almost constant thrust direction throughout the burn. The introduction of lunisolar perturbations does not alter this mechanism in any significant way, because the perturbing accelerations vary slowly over the short maneuver interval and therefore have only a limited influence on the evolution of the adjoint variables.

As a direct consequence:

- the in-plane thrust angle δ remains nearly constant and close to 90° , indicating an almost purely tangential thrust direction;

- the out-of-plane component γ is essentially null, confirming that the maneuver is performed in-plane;
- the thrust–velocity angle remains close to zero, showing that the thrust vector is nearly aligned with the instantaneous velocity direction.

This behaviour is consistent with the limited variation of λ_V visible in Fig. 6.16.1b. Since the optimal pointing law is governed by the instantaneous direction of the primer vector, the near-constancy of λ_u , λ_v , and λ_w implies a nearly time-invariant thrust direction throughout the burn arc. The thrust vector therefore remains predominantly tangential, exactly as in the unperturbed chemical-propulsion solution, reflecting the fact that the maneuver primarily aims at increasing the orbital energy in order to produce the required radial displacement at the collision epoch.

It is important to note that, although the inclusion of lunisolar perturbations introduces nonzero angular costates λ_ϑ and λ_φ , these variables do not directly determine the thrust orientation. The thrust direction is controlled exclusively by the velocity costates through the primer vector λ_V . Therefore, even if the angular part of the adjoint vector changes due to the loss of rotational symmetry in the perturbed dynamics, the thrust steering law remains essentially unaffected as long as the velocity costates exhibit the same qualitative behaviour observed in the unperturbed case.

Overall, the comparison confirms that the presence of lunisolar perturbations does not significantly modify the local structure of the optimal thrust arc for this short chemical maneuver. The velocity costates remain nearly constant, the thrust direction is almost time-invariant, and the maneuver continues to be characterized by an almost purely tangential in-plane acceleration. The main differences introduced by the perturbations therefore appear primarily in the global adjoint structure (through the angular costates) rather than in the local thrust pointing law governing the active arc.

6.2 Lunisolar Perturbations and Earth Non-Sphericity Case

6.2.1 Backward Propagation

When Earth non-sphericity is included together with lunisolar perturbations, the initial state obtained by backward propagation differs more significantly from those obtained in the previous dynamical models. The non-spherical gravity field introduces latitude-dependent accelerations and produces additional periodic variations of the orbital elements. As a consequence, the instantaneous motion is no longer consistent with a locally circular approximation even when the mean orbit remains close to circular. In particular, the orbital radius exhibits small oscillations along the orbit, which naturally lead to a nonzero radial velocity component u at generic epochs.

Backward propagation with the full perturbation model yields the following initial

state in ZEN coordinates:

$$\begin{aligned}
 r_0 &= 6969372.623914824 \text{ m}, \\
 \vartheta_0 &= -0.001375182534786 \text{ rad}, \\
 \varphi_0 &= 0.522997967582122 \text{ rad}, \\
 u_0 &= 8.493616470722914 \text{ m/s}, \\
 v_0 &= 7564.699213059743 \text{ m/s}, \\
 w_0 &= 12.238824625484511 \text{ m/s}.
 \end{aligned} \tag{6.2}$$

The presence of a clearly nonzero radial velocity component u_0 is therefore physically consistent with the perturbed dynamics. The combined action of third-body perturbations and non-central gravity terms introduces radial oscillations of the trajectory, so that the instantaneous motion cannot satisfy the condition $u \equiv 0$ that characterizes a perfectly circular orbit in the two-body model. As a result, even if the mean orbit remains close to circular, the spacecraft state at a given epoch generally corresponds to a configuration in which the radial distance is either increasing or decreasing along the trajectory.

6.2.2 Intermediate Three-arc Solution

From a numerical standpoint, the inclusion of Earth non-sphericity increases the sensitivity of the OCP and makes convergence more demanding. In the previous dynamical models the optimal maneuver could be obtained by direct continuation from the unperturbed two-arc solution, because the structure of the dynamics remained relatively close to the ideal two-body case. When the non-spherical gravity field is introduced, however, the equations of motion contain additional latitude-dependent terms that modify the local acceleration field and introduce a stronger coupling in the state-costate dynamics through the additional state dependence of the Hamiltonian. As a consequence, the shooting problem becomes more sensitive to the initial costate guess and the convergence basin of the Newton correction procedure becomes significantly narrower.

In practice, a direct continuation from the unperturbed two-arc solution may fail to converge. Even small inaccuracies in the initial adjoint estimate can lead to large deviations in the propagated trajectory, preventing the solver from satisfying the terminal boundary conditions. This effect is particularly evident when the perturbations modify the local dynamical balance along the orbit, as is the case when Earth non-sphericity is included together with lunisolar perturbations.

To improve the robustness of the numerical procedure, an auxiliary optimal-control problem with a three-arc switching structure coast-thrust-coast is first solved starting from the state in (6.2). Introducing an additional coast phase before the thrust arc relaxes the structure of the control law and enlarges the set of admissible trajectories that can satisfy the terminal constraint. This effectively reduces the sensitivity of the problem and facilitates the convergence of the indirect solver by allowing the trajectory to adapt more gradually to the perturbed dynamics.

The converged intermediate solution obtained with this structure is reported in Table 6.2.

This intermediate solution is not intended to represent the final maneuver structure. Its role is purely numerical: introducing an initial coast arc relaxes the shooting problem and facilitates the identification of a costate vector that is dynamically consistent with the perturbed model. Once this intermediate solution has converged, the costate vector

Parameter	Value
Perturbation model	Lunisolar + EGM
Switching structure	coast-thrust-coast
Start time t_{start}	-72.691996914625 s
t_1	3.075158889748 s
t_2	29.328165270343 s
Final time t_f	2912.85842378563 s
$\lambda_r(t_0)$	$1.25541303 \times 10^{-4}$
$\lambda_\vartheta(t_0)$	$-2.66291041 \times 10^{-4}$
$\lambda_\varphi(t_0)$	$-2.05548180335 \times 10^{-1}$
$\lambda_u(t_0)$	$-3.711302008 \times 10^{-3}$
$\lambda_v(t_0)$	$1.33207002118 \times 10^{-1}$
$\lambda_w(t_0)$	$7.6570641871 \times 10^{-2}$
$\lambda_m(t_0)$	$9.99990528991 \times 10^{-1}$

Table 6.2: *Inclined circular case ($i = 30^\circ$) – intermediate three-arc solution used for convergence regularization, including lunisolar perturbations and Earth non-sphericity.*

obtained after propagating the initial coast arc provides a much better initial guess for the adjoint variables of the target two-arc problem. In other words, the adjoint state is evaluated at the beginning of the thrust phase of the three-arc solution and used to initialize the corresponding variables in the thrust-coast formulation. Starting from this improved estimate, the solver can then transition reliably toward the desired two-arc structure while preserving the fixed-final-time constraint and the required terminal displacement.

6.2.3 Final Two-arc Solution

Once convergence has been achieved for the intermediate three-arc solution, the optimal-control problem is reformulated with a reduced two-arc thrust-coast structure. To construct a dynamically consistent initial guess for the final two-arc problem, the state and costate vectors of the three-arc solution are propagated forward along the initial coast arc up to the first switching instant. The resulting values are then used as the initial state and initial adjoint guess for the reduced thrust-coast formulation. In this way, the initialization of the final problem is performed exactly at the beginning of the active thrust arc, ensuring consistency between the state and costate variables in the reduced formulation.

The resulting initial state for the final optimization reads:

$$\begin{aligned}
r_0 &= 6962337.343730914 \text{ m}, \\
\vartheta_0 &= -1.584555654380545, \\
\varphi_0 &= -0.010088612403734, \\
u_0 &= -0.052114734004479 \text{ m/s}, \\
v_0 &= 6560.520986655153 \text{ m/s}, \\
w_0 &= 3786.917523093110 \text{ m/s}.
\end{aligned}$$

Using the transported costates as initial guess, the OCP converges to the final two-arc thrust-coast solution:

Parameter	Value
Perturbation model	Lunisolar + EGM
Switching structure	thrust-coast
Terminal displacement	105 m (radial)
Start time t_{start}	0.0 s
Switching time t_1	26.252858285516 s
Final time t_f	2909.783264895882 s
Burn duration Δt_{burn}	26.252858285516 s
Propellant consumed Δm	4.375 g
$\lambda_r(t_0)$	1.2573476×10^{-4}
$\lambda_\vartheta(t_0)$	$3.949951447 \times 10^{-3}$
$\lambda_\varphi(t_0)$	$-3.6647859438 \times 10^{-2}$
$\lambda_u(t_0)$	$-5.4854086 \times 10^{-4}$
$\lambda_v(t_0)$	$1.33342454086 \times 10^{-1}$
$\lambda_w(t_0)$	$7.7039212296 \times 10^{-2}$
$\lambda_m(t_0)$	$9.99990528991 \times 10^{-1}$

Table 6.3: *Inclined circular case ($i = 30^\circ$) – final fixed-final-time solution including lunisolar perturbations and Earth non-sphericity*

A comparison with the previous dynamical models highlights several relevant aspects of the optimal solution. First, the angular costates λ_ϑ and λ_φ exhibit significantly larger magnitudes compared with the lunisolar-only case. This behaviour reflects the stronger dynamical coupling introduced by Earth non-sphericity. When the gravity field departs from spherical symmetry, the spacecraft acceleration depends explicitly on latitude and longitude through the geopotential harmonics. As a consequence, the sensitivity of the terminal radial constraint with respect to the angular coordinates increases, and the optimal solution reflects this effect through larger angular costates. In other words, small variations in the angular components of the state can influence the radial evolution more strongly than in the previous models, and the adjoint variables adjust accordingly.

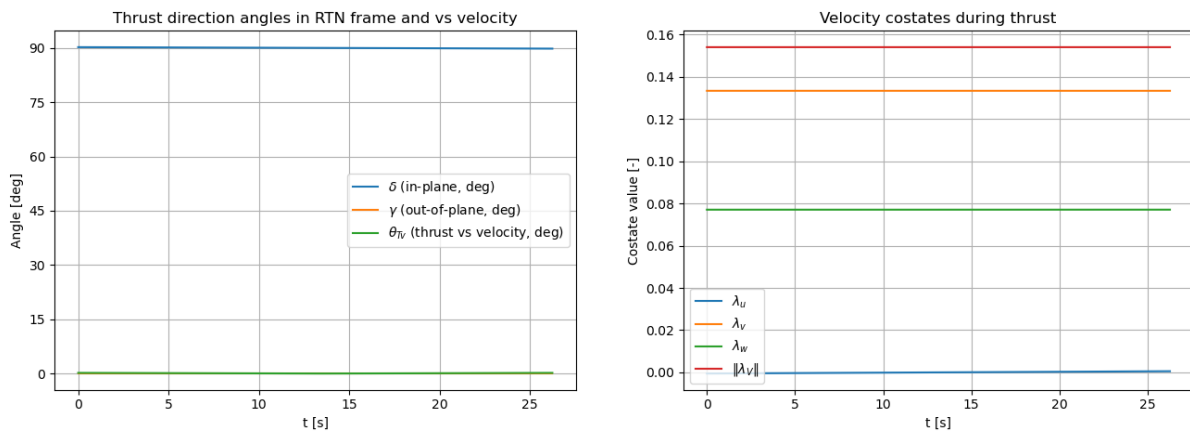
At the same time, the velocity costates λ_u , λ_v , and λ_w remain remarkably close to the values obtained in both the unperturbed and the lunisolar-only cases. Since the thrust direction is determined by the normalized velocity costate vector, this indicates that the local thrust steering law remains essentially unchanged despite the more complex perturbed dynamics. In particular, the maneuver remains governed by the same energetic mechanism identified in the previous sections: the thrust arc primarily modifies the orbital energy in order to achieve the required radial displacement at the collision epoch.

From a timing perspective, the maneuver characteristics remain almost identical to the previous chemical-propulsion solutions. The burn duration (26.25 s) is essentially the same as in both the unperturbed and the lunisolar-only cases, and the propellant consumption (4.375 g) differs only marginally from the 4.378 g obtained without perturbations. This result indicates that, over such a short maneuver interval, the additional accelerations produced by third-body perturbations and by the non-spherical gravity field generate only small corrections to the total velocity change required to achieve the prescribed 105 m radial displacement. The dominant contribution to the maneuver energetics therefore remains essentially the same as in the simpler dynamical models.

The difference in the reported final time is not associated with a different collision event. It arises from the consistent translation of the time origin together with the forward-propagated initial state used to initialize the two-arc problem. The terminal constraint is therefore still enforced at the same physical collision epoch in absolute time, while the numerical time coordinate used in the optimization is shifted consistently with the new initial state definition.

6.2.4 Thrust Direction and Velocity Costates Behavior with Lunisolar Perturbations and Earth Non-Sphericity

For completeness, the thrust-direction angles in the RTN frame and the evolution of the velocity costates during the burn arc are reported also for the case including both lunisolar perturbations and Earth non-sphericity. These quantities provide a direct visualization of the local structure of the optimal control law and allow a comparison with the solutions previously obtained in the unperturbed and lunisolar-only dynamical models.



6.2a Thrust-direction angles in the RTN frame and thrust-velocity angle during the burn.

6.2b Velocity costates and their norm during the burn.

Figure 6.2: Thrust direction and velocity costate behavior during the active thrust phase. Lunisolar perturbations and Earth non-sphericity case.

Despite the increased dynamical complexity introduced by the non-spherical gravity field, the qualitative behavior during the thrust arc remains consistent with the short-

burn solutions previously discussed. The burn duration is still of the order of 26 s, and over such a limited time interval the velocity costates λ_u , λ_v , and λ_w exhibit only very mild variations. Consequently, the norm of the primer vector $\|\lambda_V\|$ remains almost constant throughout the thrust arc, as clearly visible in Fig. 6.26.2b.

Since the optimal thrust direction is directly determined by the instantaneous direction of λ_V , the near-constancy of the velocity costates results in an almost constant thrust pointing law during the burn. In particular:

- the in-plane angle remains close to 90° , indicating a predominantly tangential thrust component, consistent with an energy-raising maneuver;
- the out-of-plane component is essentially null, confirming that the maneuver is effectively performed in-plane, in agreement with the purely radial terminal displacement requirement;
- the thrust-velocity angle shows only marginal variations over the thrust arc, indicating that the thrust vector remains almost aligned with the instantaneous velocity direction.

These characteristics are fully consistent with the behavior already observed in the unperturbed and lunisolar-only configurations. In all cases, the short duration of the thrust phase limits the time available for significant evolution of the velocity-related adjoint variables, so that the direction of the primer vector — and therefore the thrust direction — remains nearly constant during the active arc.

Overall, the results confirm that the inclusion of lunisolar perturbations and Earth non-sphericity modifies the adjoint solution at the global level but does not substantially alter the local steering law of the short chemical maneuver. The thrust arc remains nearly tangential and characterized by an almost constant pointing direction, while the perturbations primarily affect the optimal costate structure and the detailed adjustment of the trajectory required to satisfy the terminal radial displacement at the prescribed collision epoch.

6.3 Comparative Discussion

The progressive inclusion of perturbations highlights two main effects that significantly modify both the dynamical interpretation of the problem and the numerical properties of the optimal-control formulation. While the structure of the optimal maneuver remains qualitatively similar across the considered dynamical models, the underlying orbital dynamics and the sensitivity of the OCP exhibit important differences that must be carefully analyzed.

6.3.1 Loss of Orbit Equivalence and Role of Absolute Epoch

In the unperturbed two-body circular orbit, choosing a different point along the orbit as the maneuver initial condition leads to an equivalent dynamical configuration. The velocity magnitude is entirely determined by the orbital radius, and the system is autonomous: the equations of motion do not depend explicitly on time. As a consequence,

any shift of the initial condition along the circular trajectory simply corresponds to a rigid rotation of the same orbit, without modifying the intrinsic dynamical properties of the problem. The evolution of the system over a fixed time horizon is therefore invariant under a change of the initial anomaly, provided that the orbital radius remains unchanged.

With perturbations, this equivalence is lost. The initial condition depends not only on the geometric location along the orbit, but also on the absolute epoch and on the specific orbital phase at which the maneuver begins. In particular:

- third-body accelerations vary explicitly with time through the ephemerides of the perturbing bodies;
- non-spherical gravity depends on the spacecraft latitude and longitude with respect to the Earth-fixed gravity field, which rotates relative to the inertial frame;
- the instantaneous radius and velocity components can vary even for trajectories that remain close to the nominal orbit, because the orbit is no longer exactly circular and the system is no longer autonomous.

As a result, two states located at different orbital phases are no longer dynamically equivalent, even if they belong to the same nominal orbital configuration. The perturbing accelerations depend on the absolute position of the spacecraft with respect to the Earth and to the third bodies at that specific epoch. Consequently, the subsequent evolution over a fixed time horizon is strongly influenced by the exact starting point along the orbit.

For this reason, the backward propagation step becomes essential not only to ensure feasibility of the maneuver horizon, but also to generate a dynamically consistent state at the correct absolute epoch. By propagating the reference orbit backward in time while preserving absolute-time consistency in the perturbation model, a new initial state is obtained that correctly reflects the perturbed dynamics at that epoch.

In this framework, the optimization of the maneuver start time t_{start} plays a crucial role. Since the dynamics depend explicitly on time and position, the optimizer must identify the specific orbital phase that allows the prescribed terminal displacement to be achieved exactly at the required collision epoch. The combination of backward propagation and t_{start} optimization therefore enables the solver to select the appropriate point along the perturbed orbit from which the thrust arc produces the desired radial effect at the imposed final time.

In contrast to the unperturbed case, where shifting the initial anomaly does not alter the structure of the solution, the perturbed problem requires a careful synchronization between:

- the absolute epoch;
- the spacecraft position relative to Earth and third bodies;
- and the placement of the thrust window within the fixed terminal horizon.

This loss of orbit equivalence explains why the perturbed solutions exhibit modified costate structures and slightly different optimal timings, even though the terminal constraint remains purely radial. The optimal-control solution must now compensate not only for the desired energy variation, but also for the time-dependent deviations induced by the perturbations along the orbit.

6.3.2 Non-circularity and Appearance of $u \neq 0$

The presence of a nonzero radial velocity component in the lunisolar plus Earth-non-sphericity case is physically expected. In the ideal two-body circular model, the radial velocity is identically zero because the orbital radius is constant along the trajectory and the motion is purely tangential. The dynamical equilibrium between centrifugal acceleration and central gravity enforces a strictly circular trajectory with $u = \dot{r} = 0$ at all times.

Once perturbations are introduced, this property is no longer preserved. Third-body accelerations and non-spherical gravity components generate periodic and secular variations in the orbital elements. As a consequence, the radial distance is no longer constant, and small oscillations around the nominal circular radius appear. These variations naturally induce a nonzero radial velocity component, even if the orbit remains close to circular in a global sense.

In particular:

- lunisolar perturbations introduce time-dependent accelerations whose projection along the radial direction produces periodic variations of r ;
- Earth non-sphericity (through the zonal and tesseral harmonics) modifies the effective gravity field, leading to latitude-dependent accelerations that couple radial and angular motion;
- the combined effect breaks the exact balance that characterizes a Keplerian circular orbit.

Therefore, the appearance of $u \neq 0$ in the propagated initial state is not an artifact of the numerical procedure, but a direct dynamical consequence of the higher-fidelity force model. The orbit is no longer strictly circular in the Keplerian sense; instead, it exhibits small but dynamically consistent deviations in both radius and velocity components.

Although the magnitude of u remains small compared to the tangential velocity components, its presence is significant from an optimal-control perspective. The radial dynamics now evolve under a non-autonomous forcing term, and the terminal radial constraint must be satisfied on top of these intrinsic oscillations. This contributes to the increased sensitivity of the adjoint variables observed in the fully perturbed case and partially explains the stronger coupling between radial and angular costates discussed in the previous sections.

6.3.3 Increased Sensitivity and Convergence Implications

From a numerical perspective, the combined perturbation model significantly increases the sensitivity of the OCP. In the unperturbed two-body problem, the dynamics are autonomous and exhibit strong structural symmetry. Small variations in the initial adjoint guess typically produce smooth variations of the terminal residuals, and the differential-correction scheme converges robustly once a reasonable initial solution is available.

When lunisolar perturbations and Earth non-sphericity are included, the situation changes qualitatively:

- the dynamics becomes explicitly time dependent through third-body ephemerides;
- the gravity field loses spherical symmetry, introducing coupling between radial and angular motion;
- the state evolution becomes more sensitive to the precise orbital phase and absolute epoch.

As a consequence, the mapping between initial costates and terminal constraints becomes more nonlinear and less smooth. The adjoint equations are now driven by a richer and more strongly coupled Jacobian structure, which amplifies the sensitivity of the terminal conditions to small variations in the initial guess. This manifests as reduced robustness of the Newton-type differential correction and a narrower basin of convergence.

The introduction of an intermediate three-arc coast-thrust-coast solution provides a practical continuation strategy. By temporarily relaxing the structural rigidity of the two-arc solution, the optimizer is allowed to adjust the switching times more freely and to accommodate the additional coupling induced by the perturbations. Once convergence is achieved in this more flexible configuration, the resulting state and adjoint variables constitute a dynamically consistent and sufficiently accurate guess for the final two-arc thrust-coast problem.

This behaviour highlights an important methodological point: increasing model fidelity does not only modify the optimal trajectory and the adjoint structure, but also directly affects the conditioning of the indirect optimal-control problem. Higher-fidelity force models lead to stronger state-costate coupling and explicit time dependence, which in turn require more careful continuation strategies to ensure reliable convergence of the OCP.

6.4 Electric-Propulsion Cases with Lunisolar Perturbations and Earth Non-Sphericity

After obtaining a converged fixed-final-time solution including both lunisolar perturbations and Earth non-sphericity, the same perturbed optimal-control setup is solved again by modifying the propulsion parameters. This continuation step allows one to investigate how the optimal policy adapts when the thrust level and exhaust velocity

are representative of electric propulsion systems in the presence of time-dependent perturbations. The dynamical model, boundary conditions, and switching structure remain unchanged, while only the propulsion parameters (T_{\max}, c) are varied.

In this way, the influence of the propulsion system can be analyzed while preserving the same dynamical environment and terminal constraint, namely the prescribed radial displacement at the collision epoch. The perturbed chemical-propulsion solution obtained in the previous section therefore serves as a reference configuration against which the electric-propulsion cases can be directly compared.

As in the unperturbed analysis, two electric-propulsion configurations are considered:

- an ion-thruster configuration characterized by very low thrust and high exhaust velocity;
- a Hall-effect thruster configuration characterized by higher thrust and moderately high exhaust velocity.

These two propulsion models represent two typical regimes of electric propulsion and allow the investigation of how the optimal-control solution transitions from short, quasi-impulsive chemical maneuvers to distributed low-thrust control in the presence of a perturbed dynamical environment.

Compared to the perturbed chemical baseline solution (burn duration ≈ 26.25 s, propellant consumption ≈ 4.375 g), the electric-propulsion cases lead to a substantially different temporal distribution of the control. Because the available thrust level is significantly lower, the maneuver cannot deliver the required orbital-energy variation over a short interval and must instead be distributed over a longer portion of the fixed time horizon.

Consequently, both electric configurations generally exhibit:

- a substantially longer active thrust interval within the same fixed time horizon;
- a reduction in propellant consumption when the exhaust velocity is sufficiently high;
- a more pronounced time evolution of the velocity costates during the burn arc, which translates into a more gradual variation of the thrust direction in accordance with the optimal pointing law discussed in Section 4.2.3.

This behavior reflects the fundamental trade-off between thrust authority and exhaust velocity that characterizes electric propulsion. While the chemical solution concentrates the maneuver into a short interval with nearly constant thrust direction, the electric-propulsion solutions distribute the control action over a longer time span, allowing the adjoint variables associated with the velocity components to evolve more significantly during the burn.

6.4.1 Hall-Effect Thruster Case: $T_{\max} = 0.05 \text{ N}$, $c = 20000 \text{ m/s}$

The first electric-propulsion configuration considered in the perturbed dynamical environment corresponds to a propulsion system representative of Hall-effect thrusters. In this case the available thrust is increased relative to very-low-thrust electric propulsion, while the exhaust velocity remains representative of typical Hall-effect propulsion performance. This configuration therefore occupies an intermediate regime between the short chemical maneuver previously analyzed and the extremely low-thrust ion-propulsion case discussed later.

Because the thrust level is significantly lower than in the chemical baseline, the maneuver can no longer deliver the required orbital-energy variation over a short interval. Instead, the thrust arc must extend over a much longer portion of the fixed time horizon in order to achieve the prescribed radial displacement at the collision epoch. At the same time, the relatively high exhaust velocity leads to a substantial reduction in propellant consumption with respect to the chemical solution.

Parameter	Value
Propulsion system	Hall-effect thruster
T_{\max}	0.05 N
c	20000 m/s
Initial mass m_0	462.0 kg
Switching structure	thrust-coast
Terminal displacement	105 m
Final time t_f	2909.783264895882 s
Start time t_{start}	-75.577041532458 s
Switching time t_1	187.464570103223 s
Burn duration Δt_{burn}	327.003688 s
Propellant consumed Δm	0.658 g
$\lambda_r(t_0)$	1.8979301×10^{-5}
$\lambda_\vartheta(t_0)$	$-3.8661665 \times 10^{-5}$
$\lambda_\varphi(t_0)$	$-3.1011092627 \times 10^{-2}$
$\lambda_u(t_0)$	$-5.59883874 \times 10^{-4}$
$\lambda_v(t_0)$	$2.0138160847 \times 10^{-2}$
$\lambda_w(t_0)$	$1.1576154184 \times 10^{-2}$
$\lambda_m(t_0)$	$9.9999856819 \times 10^{-1}$


Table 6.4: *Inclined circular case ($i = 30^\circ$) – fixed-final-time solution with Hall-effect propulsion including perturbations.*

Compared with the perturbed chemical baseline solution, the maneuver duration increases dramatically. The burn interval grows from approximately 26 s to about 327 s, reflecting the much lower thrust authority available in this electric-propulsion regime. At the same time, the propellant consumption decreases from 4.375 g to 0.658 g thanks to the substantially higher exhaust velocity. This illustrates the classical trade-off between thrust magnitude and propellant efficiency: the maneuver becomes longer in time but significantly more efficient in terms of mass expenditure.

The longer thrust interval also affects the behavior of the adjoint variables associated with the velocity components. In contrast with the short-burn chemical case, where the velocity costates remain nearly constant during the active arc, the extended burn duration allows λ_u , λ_v , and λ_w to evolve more visibly over time. As a consequence, the direction of the primer vector varies gradually during the thrust phase.

This behavior is directly reflected in the thrust-direction angles shown in Fig. 6.3. The thrust direction is no longer almost constant as in the short-burn chemical solution, but instead exhibits a smooth reorientation over the course of the burn. Nevertheless, the maneuver still maintains a predominantly tangential character, consistent with the energy-raising mechanism required to generate the prescribed radial displacement at the collision epoch.

Although the presence of lunisolar perturbations and Earth non-sphericity modifies the global adjoint structure of the problem, the qualitative steering strategy remains similar to the unperturbed electric-propulsion cases. The thrust arc distributes the control action over a longer time interval, and the gradual evolution of the velocity costates drives a progressive adjustment of the thrust direction throughout the maneuver.



6.3a Thrust-direction angles in RTN and thrust-velocity angle during the burn.

6.3b Velocity costates and their norm during the burn.

Figure 6.3: Thrust direction and velocity costate behavior during the active thrust phase. Electric-propulsion case with $T_{\max} = 0.05$ N and $c = 20000$ m/s under lunisolar perturbations and Earth non-sphericity.

6.4.2 Ion Thruster Case: $T_{\max} = 0.01$ N, $c = 30000$ m/s

The second electric-propulsion configuration corresponds to a propulsion system representative of ion thrusters. In this case the maximum thrust is strongly reduced with respect to the Hall-effect thruster configuration, while the exhaust velocity is increased to a value typical of ion propulsion systems. This regime therefore represents the extreme low-thrust limit among the propulsion models considered in this work.

Because the available thrust is significantly smaller than in both the chemical and Hall-effect propulsion cases, the maneuver must be distributed over a much longer portion of the fixed time horizon in order to generate the required orbital-energy variation associated with the prescribed radial displacement at the collision epoch. Since the final time remains fixed at the physical collision epoch, the optimizer adapts by redistributing the thrust arc within the available horizon, resulting in an earlier optimal start time and a very long thrust interval.

Parameter	Value
Propulsion system	Ion thruster
T_{\max}	0.01 N
c	30000 m/s
Initial mass m_0	462.0 kg
Switching structure	thrust-coast
Terminal displacement	105 m
Final time t_f	2909.783264895882 s
Start time t_{start}	-73.751212620018 s
Switching time t_1	1478.311917764332 s
Burn duration Δt_{burn}	1546.872283 s
Propellant consumed Δm	0.517 g
$\lambda_r(t_0)$	2.2826238×10^{-5}
$\lambda_\theta(t_0)$	$-2.8113516 \times 10^{-5}$
$\lambda_\phi(t_0)$	$-3.6560391911 \times 10^{-2}$
$\lambda_u(t_0)$	$-6.59634203 \times 10^{-4}$
$\lambda_v(t_0)$	$2.4219544892 \times 10^{-2}$
$\lambda_w(t_0)$	$1.3925106124 \times 10^{-2}$
$\lambda_m(t_0)$	$9.99998278078 \times 10^{-1}$

Table 6.5: *Inclined circular case ($i = 30^\circ$) – fixed-final-time solution with ion propulsion including perturbations.*

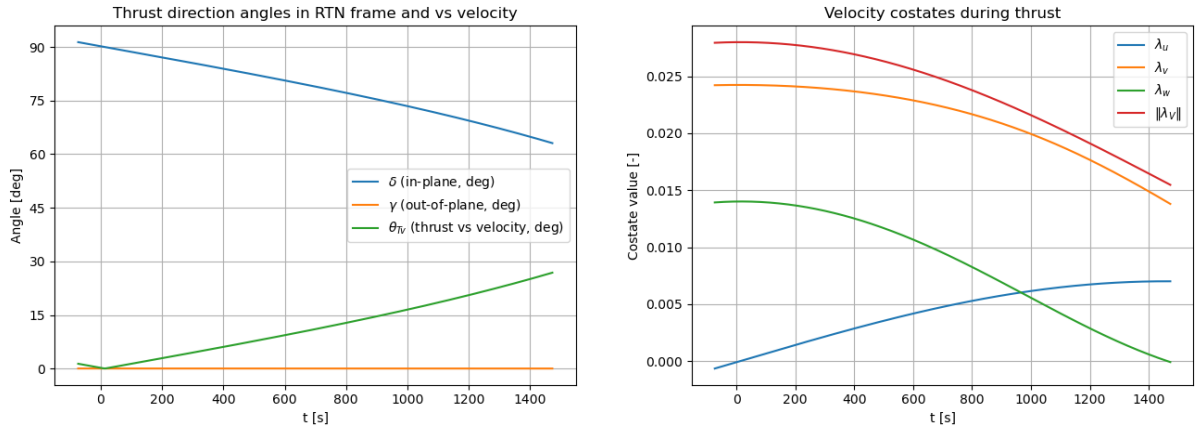
A key qualitative difference with respect to the perturbed chemical baseline solution is the dramatic increase in burn duration, which grows from approximately 26 s to more than 1500 s. This behavior is a direct consequence of the much smaller available thrust: the required variation of orbital energy must be accumulated gradually over a large portion of the available maneuver horizon. The optimizer therefore spreads the control action over a long interval in order to satisfy the terminal radial displacement at the prescribed final time.

At the same time, the propellant consumption decreases significantly, from 4.375 g in the chemical case to 0.517 g in the ion-thruster configuration. This reduction is enabled by the much higher exhaust velocity, which increases the propulsive efficiency of the maneuver. The classical trade-off between thrust authority and propellant efficiency characteristic of electric propulsion therefore remains clearly visible even when the dynamical model includes both time-dependent third-body perturbations and Earth non-sphericity.

Compared with the Hall-effect thruster case discussed in the previous subsection, the ion configuration exhibits a substantially longer thrust arc due to the lower thrust magnitude. While the Hall-effect maneuver distributes the thrust over several hundreds of seconds, the ion-propulsion solution extends the active control interval over more than half of the available time horizon. This difference directly reflects the lower acceleration capability of ion propulsion and represents the limiting low-thrust regime among the propulsion systems considered in this analysis.

Over such an extended thrust interval, the velocity costates λ_u , λ_v , and λ_w exhibit a much more pronounced time evolution than in both the chemical and Hall-effect propulsion solutions. Since the optimal thrust direction is determined by the instantaneous direction of the primer vector λ_V , the gradual evolution of these adjoint variables leads to a progressive rotation of the thrust vector during the maneuver. Consequently, the thrust direction varies continuously throughout the burn, reflecting the distributed

nature of the low-thrust control strategy.



6.4a Thrust-direction angles in RTN and thrust-velocity angle during the burn.

6.4b Velocity costates and their norm during the burn.

Figure 6.4: Thrust direction and velocity costate behavior during the active thrust phase. Electric-propulsion case with $T_{\max} = 0.01$ N and $c = 30000$ m/s under lunisolar perturbations and Earth non-sphericity.

6.4.3 Comparative Discussion

The comparison confirms that the classical propulsion trade-offs remain valid even when the dynamics includes time-dependent lunisolar perturbations and Earth non-sphericity. In particular, varying (T_{\max}, c) does not alter the overall structure of the fixed-final-time solution—which remains characterized by a bang-bang control with a single switching point—but it substantially affects the temporal distribution of the thrust arc, the associated propellant expenditure, and the evolution of the optimal thrust direction through the behavior of the velocity costates.

The perturbed chemical baseline solution provides the reference limiting case in which the required terminal radial displacement is achieved through a short, concentrated thrust arc (about 26 s). Over such a limited time interval, the velocity costates remain nearly constant and the thrust pointing law reconstructed from λ_V appears almost time-invariant. This behavior is consistent with a maneuver that is effectively impulsive with respect to the orbital time scale: the control acts over a short interval, so the state is modified mainly through the integrated effect of a nearly fixed thrust direction. In addition, the relatively low exhaust velocity leads to a higher propellant consumption ($\Delta m \approx 4.375$ g) compared to the electric-propulsion cases.

Transitioning to electric propulsion in the same perturbed dynamical environment fundamentally changes the way the terminal condition is met. Both the Hall-effect and ion-thruster configurations exhibit much longer thrust intervals because of their lower thrust authority: the same terminal radial displacement must be accumulated gradually rather than delivered over a short arc. As the active-control interval increases, the role of the velocity costates becomes more evident. Since the thrust direction is

determined instantaneously by the primer vector, a longer burn implies that the state is continuously driven by the time history of λ_V throughout the maneuver. Even moderate variations of λ_u , λ_v , and λ_w over hundreds or thousands of seconds therefore translate into a progressive reorientation of the thrust direction, which is clearly visible in the thrust-angle histories. In this sense, the differences with respect to the chemical baseline concern primarily the temporal distribution of the control and the time evolution of the thrust direction, rather than the fundamental structure of the optimal policy.

Comparing the two electric configurations, the Hall-effect thruster case represents an intermediate regime. The thrust arc extends to several minutes ($\Delta t_{\text{burn}} \approx 327$ s), which is long enough for the velocity costates to exhibit a visible evolution, yet still short compared to the ion-thruster case. The exhaust velocity remains high, yielding a substantial reduction in propellant consumption with respect to the chemical baseline ($\Delta m \approx 0.658$ g), while maintaining a burn duration that is compatible with a moderately distributed low-thrust strategy.

The ion-thruster case represents the extreme low-thrust limit. The thrust arc occupies a large fraction of the available horizon ($\Delta t_{\text{burn}} \approx 1547$ s), and the maneuver becomes strongly distributed in time. In this regime, the continuous application of the control enhances the cumulative effect of the time evolution of λ_V on the trajectory: the thrust direction must adapt progressively as the maneuver unfolds, and the thrust-angle evolution becomes correspondingly more pronounced. At the same time, the very high exhaust velocity yields the lowest propellant consumption among the considered cases ($\Delta m \approx 0.517$ g).

In all configurations, the optimization of the start time t_{start} plays a central role. Because the perturbed dynamics depends explicitly on the absolute epoch and on the instantaneous orbital configuration relative to the perturbing bodies and to the non-spherical gravity field, shifting the thrust window changes the accumulated effect of both the natural perturbations and the controlled acceleration over the remaining horizon. The solver therefore selects the maneuver placement that ensures the prescribed terminal radial displacement is met exactly at the collision epoch. Importantly, this placement is not dictated by the perturbations alone: it also depends on the propulsion parameters. Different values of T_{max} and c change the achievable acceleration profile and the duration over which thrust must be applied, thereby modifying which orbital phase yields the most effective contribution toward the terminal constraint. Consequently, the optimal t_{start} reflects the combined influence of the time-dependent perturbed environment and the specific thrust authority of the adopted propulsion system.

Overall, the perturbed comparison highlights a robust qualitative picture: the chemical case approaches an impulsive limit with almost fixed thrust direction and higher propellant cost, while the electric-propulsion cases approach progressively more distributed low-thrust regimes in which the longer active-control interval makes the evolution of the velocity costates and the associated thrust-direction adjustment a defining feature of the optimal solution.

7 Conclusions

This work investigated the design of fuel-optimal collision avoidance maneuvers in low Earth orbit through the formulation and numerical solution of a fixed-final-time OCP. The proposed framework combines a geometric analysis of collision risk with an indirect optimal-control approach based on Pontryagin's Maximum Principle, allowing the computation of minimum-propellant maneuvers capable of guaranteeing a prescribed radial separation at the predicted collision epoch.

The analysis was developed progressively by increasing the level of dynamical fidelity. Starting from an ideal two-body circular orbit, the optimal maneuver structure was first characterized in an autonomous dynamical environment. The model was then extended to include time-dependent perturbations due to third bodies and Earth non-sphericity, allowing the investigation of how environmental disturbances affect the optimal-control solution. Finally, the influence of propulsion technology was examined by comparing chemical propulsion with two representative electric-propulsion regimes.

The results provide insight into both the dynamical and numerical properties of the optimal-control formulation and highlight the main factors governing the design of efficient collision avoidance maneuvers. The following sections summarize the main conclusions of the study, discuss the limitations of the adopted modeling assumptions, and outline several possible extensions that could further improve the realism and applicability of the proposed methodology.

7.1 Interpretation of Results

The results obtained in this work highlight several important aspects of optimal collision avoidance maneuver design in low Earth orbit. In particular, the analysis clarifies the influence of the propulsion system, the dynamical environment, and the optimal-control formulation on the structure of the resulting maneuver.

A central element of the proposed framework is the geometric interpretation of collision risk through the construction of a collision-probability map in the encounter plane. By projecting the relative position uncertainty onto the plane orthogonal to the relative velocity at the time of closest approach, the two-dimensional probability density can be visualized and analyzed in a compact and intuitive manner. The resulting probability map provides valuable insight into the geometry of the uncertainty distribution and into the regions of the encounter plane that contribute most strongly to the collision probability.

The Pc map therefore plays a dual methodological role in the proposed framework. On the one hand, it provides a clear visualization of the uncertainty structure associated with the conjunction event, highlighting the orientation and extent of the covariance ellipse and the spatial distribution of collision risk. On the other hand, it offers a practical tool for the preliminary selection of a displacement direction capable of reducing the probability of collision. By inspecting the regions of highest risk in the encounter plane, it becomes possible to identify candidate displacement directions that move the spacecraft trajectory away from the most critical areas of the probability distribution.

The displacement selected through this geometric analysis is then used as the target condition for the optimal-control problem, which computes the minimum-propellant maneuver capable of producing the required orbital modification at the collision epoch. In this sense, the collision-probability map acts as an intermediate layer connecting the probabilistic characterization of the conjunction with the deterministic trajectory optimization problem.

At the same time, it is important to recognize that this geometric selection of the displacement represents only a first-order approximation. The Pc map is constructed under the assumption that the encounter geometry, the relative velocity direction, and the uncertainty structure remain unchanged when the maneuver is applied. In reality, once the spacecraft trajectory is modified by the avoidance maneuver, the relative geometry of the encounter may change. The direction of the relative velocity can rotate slightly, the encounter plane may shift, and the propagated uncertainty distribution may also evolve differently. As a result, the collision-probability map used to select the displacement does not exactly represent the post-maneuver encounter configuration.

Despite this approximation, the Pc map remains extremely valuable as an initial design tool. It allows the risk geometry to be understood rapidly and provides a rational basis for selecting candidate displacement directions before the optimal-control problem is solved. The optimal maneuver then refines this preliminary choice by computing the control profile required to achieve the selected displacement in a dynamically consistent manner.

Beyond the collision-risk analysis, the optimal-control results themselves highlight several important dynamical features. In all the considered configurations, including the cases with lunisolar perturbations and Earth non-sphericity, the optimal policy retains a bang–bang structure with a single switching point. This indicates that the fundamental nature of the maneuver does not change when higher-fidelity perturbation models are introduced. The control strategy remains characterized by an active thrust arc followed by a coast phase, with the thrust direction aligned with the instantaneous orientation of the primer vector.

The chemical-propulsion baseline case represents the limiting regime of short-duration maneuvers. Because the available thrust is relatively large, the required variation in orbital energy can be delivered over a short interval. As a consequence, the velocity costates remain nearly constant during the burn and the thrust direction appears almost time-invariant. From a dynamical perspective, this behaviour approaches the impulsive limit, in which the maneuver behaves as a nearly instantaneous velocity increment

applied at a specific orbital phase.

When electric propulsion is considered, the optimal maneuver changes significantly in its temporal distribution. Due to the much lower thrust levels, the required orbital-energy variation must be accumulated gradually over a longer portion of the available time horizon. The thrust arc therefore becomes substantially longer, extending from several minutes in the Hall-effect thruster case to more than half of the available maneuver horizon in the ion-thruster configuration. This longer active interval allows the velocity costates to evolve more noticeably over time, producing a gradual reorientation of the thrust direction during the maneuver. The thrust steering law therefore reflects the distributed nature of the low-thrust control strategy.

Another important result concerns the classical trade-off between thrust authority and propellant efficiency. The chemical solution achieves the required orbital modification with the shortest maneuver duration but exhibits the largest propellant consumption. In contrast, electric propulsion significantly reduces the propellant mass required to achieve the same terminal condition, thanks to the higher exhaust velocity. The ion-thruster configuration achieves the lowest propellant consumption but requires the longest maneuver duration, while the Hall-effect thruster represents an intermediate regime between these two extremes.

The introduction of perturbations provides additional insight into the dynamical structure of the problem. In the unperturbed two-body model, the system is autonomous and circular two-body orbits exhibit rotational symmetry. As a consequence, the choice of the initial orbital position does not affect the intrinsic dynamical properties of the maneuver. When perturbations are included, this symmetry is lost. Third-body accelerations depend explicitly on time through the ephemerides of the perturbing bodies, while Earth non-sphericity introduces latitude- and longitude-dependent accelerations through the geopotential harmonics. The optimal maneuver therefore depends on the absolute epoch and on the instantaneous orbital configuration with respect to the perturbing bodies.

This loss of orbital equivalence also affects the optimal-control problem from a numerical perspective. The mapping between initial costates and terminal constraints becomes more sensitive when perturbations are included, reducing the robustness of the differential-correction procedure used to solve the OCP. The introduction of an intermediate three-arc continuation strategy proved essential for achieving reliable convergence in the fully perturbed dynamical model.

Despite these dynamical and numerical differences, the overall structure of the optimal maneuver remains consistent across all scenarios considered in this work. In particular, the thrust arc always acts primarily as an energy-modifying maneuver, producing the required radial displacement at the collision epoch while preserving the overall geometry of the orbit. This indicates that the proposed optimal-control formulation provides a robust framework for the design of fuel-efficient collision avoidance maneuvers even in the presence of realistic perturbations.

7.2 Model Limitations

Despite the encouraging results obtained in this study, several simplifying assumptions were introduced in order to keep the problem tractable and to focus on the fundamental structure of the optimal-control formulation. These assumptions inevitably limit the physical realism of the model and suggest several directions for further refinement.

A first limitation concerns the representation of the dynamical environment. Although the final model includes both lunisolar perturbations and Earth non-sphericity through the EGM2008 geopotential, other perturbative effects that can be relevant in low Earth orbit have not been considered. In particular, atmospheric drag and solar radiation pressure were neglected. While these perturbations are generally smaller than the dominant gravitational terms over the relatively short time horizon considered in this work, they may still produce measurable deviations in the orbital evolution, especially for satellites with large area-to-mass ratios or at lower altitudes. Including these additional forces would increase the realism of the trajectory propagation and could influence the optimal maneuver timing and control profile.

A second limitation concerns the accuracy of the geopotential representation itself. In the present analysis the non-spherical gravity field is introduced through a truncated spherical harmonic expansion, assuming a simplified Earth orientation model. In reality, a fully consistent high-fidelity gravity model would also require accounting for the time-dependent orientation of the Earth, including precession, nutation, and tidal effects. These phenomena modify the orientation of the geopotential coefficients in inertial space and therefore influence the long-term orbital evolution. Although their impact is typically small over the short maneuver horizons considered here, their inclusion would be necessary for a fully rigorous high-fidelity dynamical model.

Another important limitation concerns the modeling of the collision risk itself. In the present work the covariance associated with the encounter is not obtained from real tracking data but is instead estimated statistically in order to construct a representative uncertainty distribution in the encounter plane. This approach allows the geometric structure of the collision probability map to be studied and used as a basis for maneuver design, but it does not reproduce the full complexity of real conjunction scenarios. In operational environments, covariance matrices are derived from orbit determination processes based on tracking measurements, and their evolution depends on both the dynamical model and the observation geometry.

The use of a statistically generated covariance introduces a further limitation when the maneuver is applied. Because the covariance is defined at the predicted collision epoch and is not propagated consistently from the maneuver start time, the uncertainty distribution used in the collision analysis is not dynamically consistent with the trajectory modification produced by the optimal maneuver. As a consequence, it is not possible within the current framework to perform a fully coherent post-maneuver collision probability assessment starting from the maneuver epoch.

A further simplification concerns the geometric modeling of the conjunction event. The collision probability computation relies on the classical short-encounter approximation,

which reduces the three-dimensional relative-motion problem to a two-dimensional formulation in the encounter plane. This approximation is widely used in operational conjunction analysis and is generally valid for high-velocity encounters, but it neglects the full three-dimensional structure of the uncertainty distribution and the detailed time evolution of the relative trajectory during the encounter.

An additional limitation is related to the use of the collision-probability map in the encounter plane to guide the selection of the avoidance displacement. The Pc map is constructed assuming the nominal encounter geometry defined by the relative velocity at the predicted time of closest approach. The encounter plane is therefore oriented according to this relative velocity vector, and the uncertainty distribution is projected onto that plane to identify the regions that contribute most strongly to the collision probability. The displacement direction is then selected within this fixed geometric configuration.

However, once an avoidance maneuver is applied, the relative geometry of the encounter may change. The spacecraft trajectory is modified, which may lead to variations in the relative velocity vector and consequently in the orientation of the encounter plane itself. In addition, the propagated uncertainty distribution associated with the modified trajectory may evolve differently. As a result, the encounter-plane configuration used to construct the Pc map does not necessarily coincide with the actual geometry of the encounter after the maneuver has been performed.

The selection of the displacement in the encounter plane should therefore be interpreted as a first-order approximation. The Pc map provides a valuable tool to identify regions of higher risk and to determine a displacement direction that moves the spacecraft away from the most critical portions of the uncertainty distribution. Nevertheless, the approach implicitly assumes that the encounter-plane geometry remains approximately unchanged after the maneuver, which is not strictly true in a fully consistent dynamical sense.

Finally, the formulation of the optimal-control problem itself includes simplifying assumptions regarding the terminal constraints. In the present work the maneuver objective is expressed through a prescribed radial displacement at the collision epoch. While this condition is sufficient to demonstrate the effectiveness of the optimal-control approach and to generate meaningful avoidance strategies, it represents only one possible formulation of the collision avoidance requirement. Real mission scenarios often involve more complex constraints on the final orbital state or on the relative geometry between the spacecraft and the debris object.

7.3 Possible Extensions

The limitations discussed above naturally suggest several possible extensions that could significantly enhance the realism and applicability of the proposed framework.

A first line of improvement concerns the dynamical model. Future developments could incorporate additional perturbations such as atmospheric drag and solar radiation pressure, which can play a non-negligible role in low Earth orbit, especially for

satellites operating at lower altitudes or with large surface-to-mass ratios. Furthermore, a more complete treatment of Earth non-sphericity could be implemented by including higher-order geopotential harmonics together with the effects of Earth orientation variations, such as precession, nutation, and tidal deformations. Although these effects are relatively small over short time intervals, their inclusion would allow the optimal-control framework to operate within a fully high-fidelity dynamical environment.

A second important extension concerns the treatment of uncertainty and collision probability. A more realistic approach would involve the use of covariance matrices obtained from actual orbit determination processes rather than statistically generated distributions. In operational conjunction assessment, these covariances are derived from tracking measurements and propagated forward in time using the same dynamical model adopted for the nominal trajectory. Integrating such covariance information into the present framework would allow the uncertainty evolution to be treated consistently throughout the maneuver planning process.

Using realistic covariance propagation would also enable a fully consistent post-maneuver collision risk assessment. Once an optimal avoidance trajectory is computed, the covariance could be propagated from the maneuver epoch to the conjunction time and the resulting probability of collision could be evaluated directly for the modified trajectory. This capability would allow the effectiveness of the maneuver to be verified quantitatively and would represent a crucial step toward operational applicability.

In addition, the availability of dynamically propagated covariance information would open the possibility of implementing an iterative maneuver-design procedure. In the current framework, the collision-probability map in the encounter plane is used only once to identify a suitable displacement direction before solving the optimal-control problem. However, once the avoidance maneuver is computed, the modified trajectory may lead to a different relative encounter geometry, potentially changing both the orientation of the encounter plane and the structure of the uncertainty distribution.

With realistic covariance propagation, it would be possible to recompute the collision probability after the maneuver by propagating both the state and the covariance from the maneuver epoch to the conjunction time. A new collision-probability map could then be generated using the updated encounter geometry. If the resulting probability distribution differs significantly from the original one, the displacement direction could be updated and a new optimal-control problem solved.

Such an iterative loop — consisting of displacement selection in the encounter plane, optimal maneuver computation, trajectory propagation, and updated collision-probability evaluation — would progressively refine the avoidance strategy. The Pc map would therefore act as a geometric guide for selecting candidate displacement regions, while the optimal-control solver would determine the dynamically feasible maneuver required to reach them. Repeating this process until convergence would produce an avoidance maneuver that is fully consistent with both the dynamical evolution and the probabilistic characterization of the encounter.

Another possible extension concerns the geometric modeling of the conjunction event.

Instead of relying on the short-encounter approximation, the collision probability could be evaluated using fully three-dimensional models of the relative motion and uncertainty distribution. Such models would capture the full time evolution of the relative trajectory and could provide a more accurate description of the collision geometry, particularly in cases where the short-encounter assumptions are not strictly satisfied.

The optimal-control formulation itself could also be extended by introducing more general terminal conditions. In the present study the maneuver objective is defined through a prescribed radial displacement at the collision epoch, which represents a convenient proxy for collision avoidance. In more general scenarios, additional constraints could be imposed on the final state of the spacecraft. For instance, one could require the spacecraft to reach a specific orbital configuration after the maneuver, impose constraints on the final velocity components, or enforce the reconstruction of a particular operational orbit following the avoidance action. Such formulations would lead to richer optimal-control problems and could better reflect the requirements of real mission operations.

Another important extension concerns the orbital configurations considered in the analysis. The present work focuses on nearly circular low Earth orbits, which simplifies both the dynamical analysis and the interpretation of the results. However, many operational satellites operate on eccentric orbits or in more complex dynamical regimes. Extending the optimal-control framework to elliptical orbits would introduce additional challenges, including a stronger time dependence of the orbital state, variations in orbital velocity along the trajectory, and a more complex relationship between thrust application and orbital energy modification. Addressing these aspects would further broaden the applicability of the proposed methodology.

Finally, improvements could also be introduced at the numerical level. The indirect single-shooting approach adopted in this work provides high accuracy and a clear interpretation of the optimality conditions through the costate variables and the primer vector. However, indirect methods can suffer from sensitivity to the initial guess, especially in strongly perturbed environments. Future work could explore hybrid numerical strategies combining indirect methods with continuation techniques or direct optimal-control formulations, potentially improving robustness and convergence properties for more complex mission scenarios.

Overall, the framework developed in this thesis provides a flexible foundation for the optimal design of collision avoidance maneuvers. By progressively incorporating higher-fidelity dynamical models, realistic uncertainty propagation, and more general mission constraints, the methodology could evolve toward a powerful tool for supporting operational conjunction mitigation strategies in increasingly congested orbital environments.

References

- [1] ESA Space Debris Office. *ESA's Annual Space Environment Report*. Issue 9.1, Final. Document ID: GEN-DB-LOG-00288-OPS-SD. ESA ESOC, Darmstadt, Germany. Oct. 2025.
- [2] Arthur E. Bryson and Yu-Chi Ho. *Applied Optimal Control: Optimization, Estimation and Control*. Taylor & Francis Group, 1975. ISBN: 9781315137667. DOI: 10.1201/9781315137667.
- [3] John T. Betts. *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. 2nd ed. Society for Industrial and Applied Mathematics, 2010. ISBN: 9780898716887. DOI: 10.1137/1.9780898718577.
- [4] Bruce A. Conway. *Spacecraft Trajectory Optimization*. Cambridge University Press, 2010. ISBN: 9780511778025. DOI: 10.1017/CB09780511778025.
- [5] Lorenzo Casalino. “Singular Arcs During Aerocruise”. In: *Journal of Guidance, Control, and Dynamics* 23(1) (2000), pp. 118–123. DOI: 10.2514/2.4494.
- [6] R. Bertrand and R. Epenoy. “New Smoothing Techniques for Solving Bang-Bang Optimal Control Problems—Numerical Results and Statistical Interpretation”. In: *Optimal Control Applications and Methods* 23(4) (2002), pp. 171–197. DOI: 10.1002/oca.709.
- [7] Morris Handelsman. “Optimal Free-Space Fixed-Thrust Trajectories Using Impulsive Trajectories as Starting Iteratives”. In: *AIAA Journal* 4(6) (1966), pp. 1077–1082. DOI: 10.2514/3.3607.
- [8] John E. Prussing and Jenq-Hua Chiu. “Optimal Multiple-Impulse Time-Fixed Rendezvous Between Circular Orbits”. In: *Journal of Guidance, Control, and Dynamics* 9(1) (1986), pp. 17–22. DOI: 10.2514/3.20060.
- [9] Luigi Mascolo. “Low-Thrust Optimal Escape Trajectories from Lagrangian Points and Quasi-Periodic Orbits in a High-Fidelity Model”. PhD thesis. Politecnico di Torino, 2021.
- [10] Luigi Mascolo. “OCULUS – A Unified Visual Solver for Optimal Control Problems via Indirect Methods”. In: *Proceedings of the 76th International Astronautical Congress (IAC 2025)*. Paper code: IAC-25,E1,4,4,x103041. Paper ID: 103041. Sydney, Australia: International Astronautical Federation (IAF), Sept. 2025.
- [11] National Aeronautics and Space Administration. *Navigation Data – Definitions and Conventions*. Document ID: CCSDS 500.0-G-4. 2019.

- [12] National Geospatial-Intelligence Agency (NGA). *EGM2008: The Earth Gravitational Model 2008*. Tech. rep. Model coefficients and documentation, Tide-Free system. 2008. URL: <https://earth-info.nga.mil>.
- [13] N. K. Pavlis, S. A. Holmes, S. C. Kenyon, and J. K. Factor. “The development and evaluation of the Earth Gravitational Model 2008 (EGM2008)”. In: *Journal of Geophysical Research: Solid Earth* 117(B4) (2012). DOI: 10.1029/2011JB008916.
- [14] David A. Vallado. *Fundamentals of Astrodynamics and Applications*. 2nd ed. Microcosm Press, 2001. ISBN: 9781881883128.
- [15] L. Mascolo and L. Casalino. “Optimal Escape from Sun-Earth and Earth-Moon L2 with Electric Propulsion”. In: *Aerospace* 9(4) (2022). DOI: 10.3390/aerospace9040186.
- [16] S. Simeoni, L. Casalino, A. Zavoli, and G. Colasurdo. “Optimal control for low-thrust Earth escape trajectories in a high-fidelity dynamical model”. In: *International Journal of Aerospace Engineering* 20 (2012), pp. 389–404. DOI: 10.1155/2012/152683.
- [17] S. A. Holmes and W. E. Featherstone. “A unified approach to the Clenshaw summation and the recursive computation of very high degree and order normalised associated Legendre functions”. In: *Journal of Geodesy* 76 (2002), pp. 279–299. DOI: 10.1007/s00190-002-0216-2.
- [18] W. Magnus, F. Oberhettinger, and R. P. Soni. *Formulas and Theorems for the Special Functions of Mathematical Physics*. Springer, 1966. ISBN: 9783662117637. DOI: 10.1007/978-3-662-11761-3.
- [19] Maruthi R. Akella and Kyle T. Alfriend. “Probability of Collision Between Space Objects”. In: *Journal of Guidance, Control, and Dynamics* 23(5) (2000), pp. 769–772. DOI: 10.2514/2.4611.
- [20] National Aeronautics and Space Administration (NASA). *NASA Spacecraft Conjunction Assessment and Collision Avoidance Best Practices Handbook, Rev. 1: Appendix N. P_c Calculation Approaches (Update 20 March 2024)*. Office of the Chief Engineer. Appendix N update to NASA/SP-20230002470 Rev. 1. Mar. 2024.
- [21] J. Foster and H. Estes. *A Parametric Analysis of Orbital Debris Collision Probability and Maneuver Rate for Space Vehicles*. Document ID: JSC-25898. 1992.
- [22] Richard Johnson and Dean Wichern. *Applied Multivariate Statistical Analysis*. 6th ed. Pearson Prentice Hall, 2013. ISBN: 9781292024943.
- [23] Kanti V. Mardia, John T. Kent, and John M. Bibby. *Multivariate Analysis*. Academic Press, 1979. ISBN: 9780124712508.
- [24] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006. ISBN: 9780387310732.
- [25] Derek F. Lawden. *Optimal Trajectories for Space Navigation*. Butterworths, 1963.

- [26] Donald E. Kirk. *Optimal Control Theory: An Introduction*. Dover Publications, 1970. ISBN: 9780486434841.
- [27] David Burghes and Alexander Graham. *Control and Optimal Control Theories with Applications*. Horwood Publishing Limited, 2004. ISBN: 9781904275015. DOI: 10.1533/9780857099495.
- [28] National Aeronautics and Space Administration (NASA). *State-of-the-Art Small Spacecraft Technology*. Document ID: NASA/TP-20250000142. Feb. 2025.