

POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Near-Optimal Bang-Bang Control for Low-Thrust Collision Avoidance Using Differential Correction

Tesi di Laurea Magistrale in
Space Engineering - Ingegneria Spaziale

Author: **Marco Barbieri**

Student ID: 243000
Advisor: Prof. Pierluigi Di Lizia
Co-advisors: Prof. Haijun Rong
Academic Year: 2025-26

Abstract

The exponential proliferation of large satellite constellations has led to unprecedented congestion in Low Earth Orbit (LEO). As the frequency of orbital conjunctions steadily rises, traditional Space Traffic Management (STM) frameworks are becoming overwhelmed, necessitating the development of autonomous, robust, and fuel-efficient Collision Avoidance Maneuvers (CAMs). While low-thrust electric propulsion systems offer superior mass efficiency for these maneuvers, their extended operational nature severely complicates trajectory optimization. Formulating a strictly fuel-optimal CAM requires deriving a discontinuous, "bang-bang" control law. Traditionally, this is achieved by solving a highly sensitive Two-Point Boundary Value Problem (TPBVP) that frequently suffers from severe numerical instability and convergence failures. The core research, mathematical formulations, and operational evaluations of this study are structured to achieve the following primary objectives:

1) Low-thrust satellite collision avoidance modeling is developed to establish a comprehensive baseline for the fuel-optimal control design. This model rigorously defines the physical and mathematical framework for short-term encounter analysis, detailing essential coordinate transformations, decoupled low-thrust two-body dynamics, and B-plane kinematics. By propagating the samples from the European Space Agency (ESA) Collision Avoidance Challenge dataset without executing CAMs, the baseline simulation reveals a severe concentration of high-risk conjunctions. Consequently, this foundation successfully extracts and formulates collision safety constraints utilizing both the deterministic physical Miss Distance and the probabilistic Squared Mahalanobis Distance, accommodating varying operational paradigms.

2) Fuel Optimal Bang-Bang Control Using Differential Correction is proposed as a novel, highly efficient methodology to overcome profound computational bottlenecks in low-thrust CAM design. This approach bypasses the numerical integration of sensitive costate differential equations entirely. By leveraging a closed-form analytical energy-optimal solution, the algorithm freezes the near-optimal thrust pointing direction and seeds a rapid bisection method to estimate the initial engine ignition and cutoff times. An iterative differential correction loop then employs the State Transition Matrix (STM) to linearly

map these temporal perturbations directly to terminal spatial deviations on the B-plane. As a result, this two-phase architecture rapidly converges on a physically realizable, fuel-optimal bang-bang thrust profile, seamlessly bridging analytical efficiency with numerical precision.

3) Performance Evaluation of the Low-Thrust Fuel Optimal Differential Correction is conducted to rigorously validate the operational efficiency, accuracy, and propellant usage of the proposed approach. This validation numerically evaluates the algorithm against state-of-the-art analytical and numerical solvers using the exact same comprehensive ESA dataset evaluated in the baseline analysis. The results demonstrate the solver's robustness, confirming a near-perfect convergence rate, sub-second execution speeds, and minimized propellant consumption across a highly diverse spectrum of strict geometric and statistical mission safety constraints, while preserving high operational accuracy even when subjected to severe environmental perturbations.

Keywords: Space Traffic Management; Collision Avoidance Maneuver; Low-Thrust Propulsion; Differential Correction; Trajectory Optimization; Probability of Collision; B-Plane Dynamics

Abstract in lingua italiana

La proliferazione esponenziale di grandi costellazioni satellitari ha portato a una congestione senza precedenti nell'orbita terrestre bassa (LEO). Con il costante aumento della frequenza delle congiunzioni orbitali, i tradizionali sistemi di gestione del traffico spaziale (Space Traffic Management - STM) stanno diventando sovraccarichi, rendendo necessario lo sviluppo di manovre di evitamento delle collisioni (Collision Avoidance Maneuvers - CAM) che siano autonome, robuste ed efficienti in termini di consumo di propellente. Sebbene i sistemi di propulsione elettrica a bassa spinta offrano un'efficienza di massa superiore per queste manovre, la loro natura operativa prolungata complica gravemente l'ottimizzazione della traiettoria. La formulazione di una CAM strettamente ottimale in termini di propellente richiede la derivazione di una legge di controllo discontinua, di tipo "bang-bang". Tradizionalmente, ciò si ottiene risolvendo un problema ai valori al contorno a due punti (Two-Point Boundary Value Problem - TPBVP) altamente sensibile, che soffre frequentemente di grave instabilità numerica e fallimenti nella convergenza. La ricerca principale, le formulazioni matematiche e le valutazioni operative di questo studio sono strutturate per raggiungere i seguenti obiettivi primari:

1) La modellazione dell'evitamento delle collisioni satellitari a bassa spinta è sviluppata per stabilire una base di riferimento completa per il progetto di controllo ottimale in termini di propellente. Questo modello definisce rigorosamente il quadro fisico e matematico per l'analisi degli incontri a breve termine, dettagliando le trasformazioni di coordinate essenziali, la dinamica disaccoppiata dei due corpi a bassa spinta e la cinematica sul B-Plane. Propagando i campioni dal dataset della Collision Avoidance Challenge dell'Agenzia Spaziale Europea (ESA) senza eseguire le CAM, la simulazione di base rivela una grave concentrazione di congiunzioni ad alto rischio. Di conseguenza, questa base estrae e formula con successo i vincoli di sicurezza contro le collisioni utilizzando sia la distanza di mancato impatto (Miss Distance) fisica e deterministica, sia la distanza di Mahalanobis al quadrato probabilistica, adattandosi a diversi paradigmi operativi.

2) Il Controllo Bang-Bang Ottimale in termini di Propellente tramite Correzione Differenziale è proposto come una metodologia nuova e altamente efficiente per superare i profondi colli di bottiglia computazionali nella progettazione di CAM a bassa spinta. Questo ap-

proccio aggira completamente l'integrazione numerica delle sensibili equazioni differenziali dei costati. Sfruttando una soluzione analitica in forma chiusa ottimale in termini di energia, l'algoritmo congela la direzione di puntamento della spinta quasi ottimale e innesca un rapido metodo di bisezione per stimare i tempi iniziali di accensione e spegnimento del motore. Un ciclo iterativo di correzione differenziale impiega poi la Matrice di Transizione di Stato (State Transition Matrix - STM) per mappare linearmente queste perturbazioni temporali direttamente nelle deviazioni spaziali terminali sul B-Plane. Di conseguenza, questa architettura in due fasi converge rapidamente su un profilo di spinta bang-bang fisicamente realizzabile e ottimale in termini di propellente, unendo senza soluzione di continuità l'efficienza analitica con la precisione numerica.

3) La Valutazione delle Prestazioni della Correzione Differenziale Ottimale per Propellente a Bassa Spinta viene condotta per validare rigorosamente l'efficienza operativa, l'accuratezza e l'utilizzo del propellente dell'approccio proposto. Questa validazione valuta numericamente l'algoritmo rispetto ai risolutori analitici e numerici all'avanguardia utilizzando l'esatto e medesimo dataset completo dell'ESA valutato nell'analisi di base. I risultati dimostrano la robustezza del risolutore, confermando un tasso di convergenza quasi perfetto, velocità di esecuzione inferiori al secondo e un consumo di propellente minimizzato attraverso uno spettro altamente diversificato di rigidi vincoli geometrici e statistici per la sicurezza della missione, pur preservando un'elevata accuratezza operativa anche quando sottoposto a severe perturbazioni ambientali.

Parole chiave: Gestione del Traffico Spaziale; Manovra di Evitamento delle Collisioni; Propulsione a Bassa Spinta; Correzione Differenziale; Ottimizzazione della Traiettoria; Probabilità di Collisione; Dinamica del B-Plane

Contents

Abstract **i**

Abstract in lingua italiana **iii**

Contents **v**

1 Introduction **1**
1.1 Research Background and Significance 1
1.2 Research Status at Home and Abroad 4
1.2.1 Analytical and Semi-analytical Approaches 5
1.2.2 Numerical Optimization Approaches 6
1.3 Main Content and Structure of the Thesis 9

2 Low-Thrust Satellite Collision Avoidance Modelling **13**
2.1 Introduction 13
2.2 Dynamics of the Satellites: Low-Thrust Two-Body Problem 13
2.2.1 Coordinate Systems 13
2.2.2 B-plane Definition 16
2.2.3 Dynamics of the Primary Satellite 17
2.2.4 Dynamics of the Secondary Satellite 18
2.3 Evaluation of Collision Conditions 19
2.3.1 Miss Distance 19
2.3.2 Probability of Collision and Squared Mahalanobis Distance 19
2.3.3 Conjunction Data Message 23
2.4 Simulation Results: Nominal Encounters Without CAM 24
2.4.1 Construction of Avoidance Constraints 24
2.4.2 Encounter Distribution and Heatmap Generation 25
2.5 Summary of This Chapter 27

3 Fuel Optimal Bang-Bang Control Using Differential Correction 29
3.1 Introduction . . . 29
3.2 Problem Formulation . . . 29
3.3 Method Description . . . 31
3.3.1 Energy Optimal Analytical Guess . . . 31
3.3.2 Fuel Optimal Control Design . . . 36
3.4 Algorithm Workflow: Energy-Optimal Initialization and Differential Correction Loop . . . 43
3.4.1 Phase 1: Energy-Optimal Initialization . . . 43
3.4.2 Phase 2: Differential Correction Loop . . . 47
3.5 Summary of This Chapter . . . 50

4 Performance Evaluation of the Low-Thrust Fuel Optimal Differential Correction 51
4.1 Introduction . . . 51
4.2 Miss Distance Constraints Results . . . 53
4.2.1 Convergence and Accuracy . . . 53
4.2.2 Propellant Consumption: Minimum Required ΔV Analysis . . . 54
4.2.3 Optimal Maneuver Lead Time Analysis . . . 57
4.2.4 Cumulative Distribution of Constraint Violations . . . 58
4.2.5 Visualization of the First Sample: ID 1 . . . 59
4.3 Squared Mahalanobis Distance Constraints Results . . . 61
4.3.1 Convergence and Accuracy . . . 61
4.3.2 Propellant Consumption: Minimum Required ΔV Analysis . . . 62
4.3.3 Optimal Maneuver Lead Time Analysis . . . 64
4.3.4 Cumulative Distribution of Constraint Violations . . . 65
4.3.5 Visualization of the First Sample: ID 1 . . . 66
4.4 Effects of Drag and J_2 Perturbations on Method Accuracy . . . 68
4.4.1 Impact on Miss Distance Constraints . . . 70
4.4.2 Impact on Probability of Collision Constraints . . . 71
4.5 Summary of This Chapter . . . 72

5 Conclusions and Future Work 75
5.1 Future Work . . . 77

Bibliography 79

List of Figures **87**

List of Tables **91**

List of Symbols **93**

Acknowledgements **95**

1 | Introduction

1.1. Research Background and Significance

The landscape of space exploration and utilization has undergone a dramatic paradigm shift over the last decade. Historically dominated by a select few government space agencies, access to space has become increasingly democratized. The commercial space sector has seen explosive growth, spearheaded by massive private enterprises such as SpaceX (Starlink), Amazon (Project Kuiper), and various state-backed ventures deploying unprecedented numbers of satellites into Low Earth Orbit (LEO).

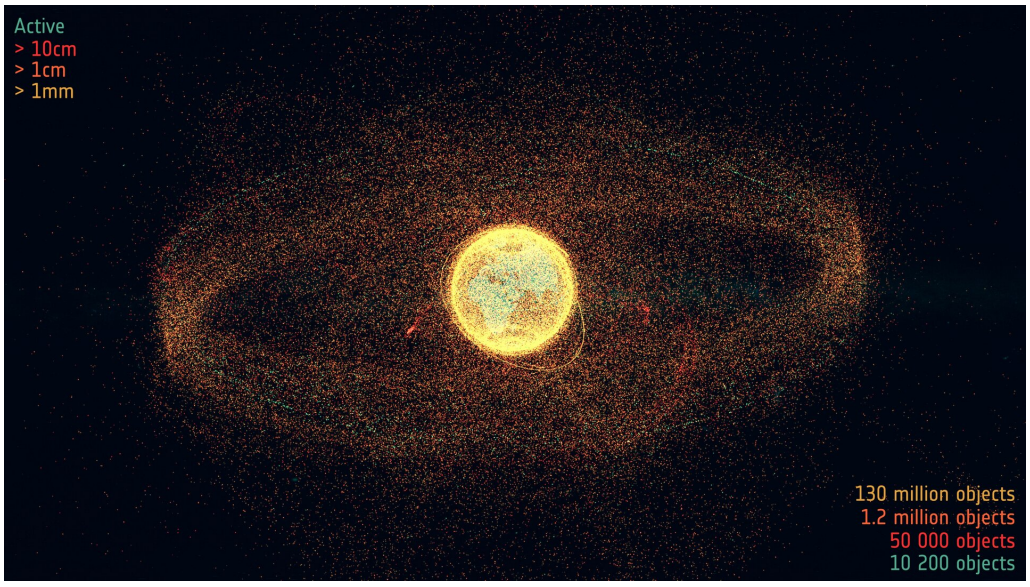

Figure 1.1: Overview of space debris distribution around the Earth [23]

This surge in accessibility translates directly to a skyrocketing launch cadence, as depicted in Figure 1.2. Consequently, the orbital population is increasing at a staggering rate, inherently increasing the probability of orbital conjunctions. Without robust mitigation strategies [4, 34], potential collisions [62] risk triggering an exponential generation of space debris—often referred to as the Kessler Syndrome, illustrated in Figures 1.4 and 1.3—which could render critical orbital regimes unusable for future generations, a

challenge further compounded by ongoing legal ambiguities in defining space debris [47].

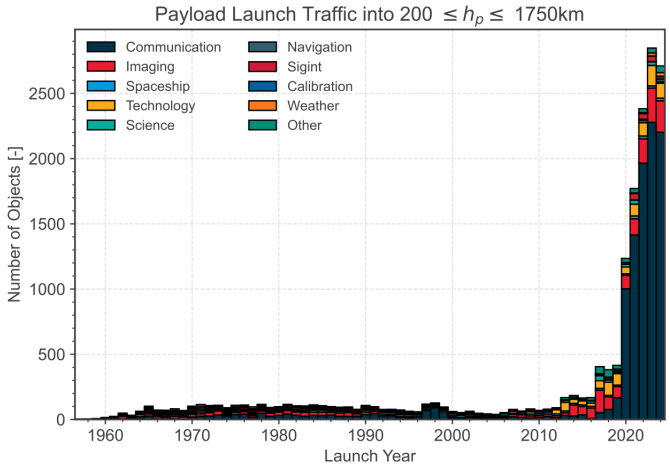

Figure 1.2: Evolution of the launch traffic near LEO per mission type in object number [23]

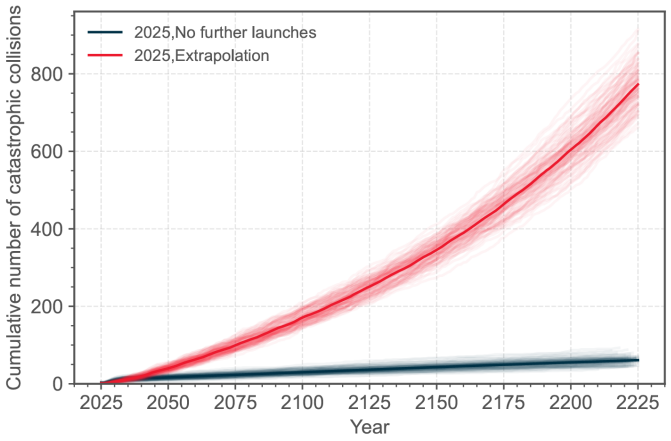

Figure 1.3: Number of cumulative collisions in LEO in the simulated scenarios of long-term evolution of the environment.

This theoretical risk has materialized into an operational reality, and can be visualized in Figure 1.1. With over 54,000 tracked objects larger than 10cm and an estimated 1.2 million lethal, untrackable debris fragments of dimensions between 1cm and 10cm [24], space surveillance networks are overwhelmed [3, 7, 38]. The United States Space Force currently generates approximately 600,000 Conjunction Data Messages (CDMs) [13] per day—an increase of 200% over the daily rate from just three years ago [32]. Consequently, mega-constellation operators and Space Traffic Management systems [53] are forced into unprecedented action; Starlink, for example, executed approximately 50,000 collision avoidance maneuvers between late 2020 and mid-2025 alone, as shown in Figure 1.5 [49].

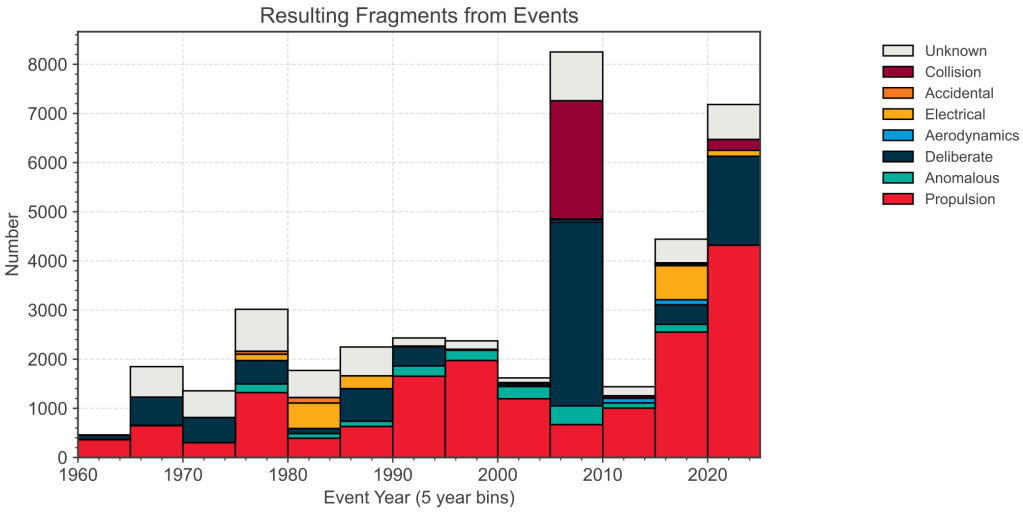

Figure 1.4: Historical trend of number of fragments produced by fragmentation events. The last bin covers until the end of 2024 [23]

To mitigate these risks, the standard approach is treated as an asymmetrical problem: one active, maneuverable spacecraft attempting to evade a passive, non-cooperative space object[28], as illustrated in Figure 1.6. Although active capture technologies like robotic nets are being researched [14], evasion remains the immediate operational necessity. This places an immense burden on the active satellite's autonomous Guidance, Navigation, and Control (GNC) subsystem [40].

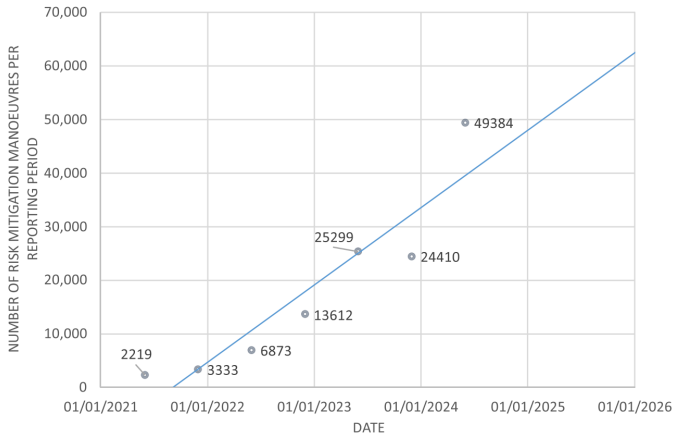

Figure 1.5: Number of conjunction risk mitigation maneuvers performed by satellites from the Starlink constellation in six-monthly periods covering December 2020 through May 2024. The linear model has a gradient of approximately 40 maneuvers per day [32]

When designing these Collision Avoidance Maneuvers (CAMs) [48], the choice of propulsion system fundamentally dictates the mathematical approach. Historically, maneuvers have been impulsive, utilizing chemical propulsion to apply high-magnitude forces over

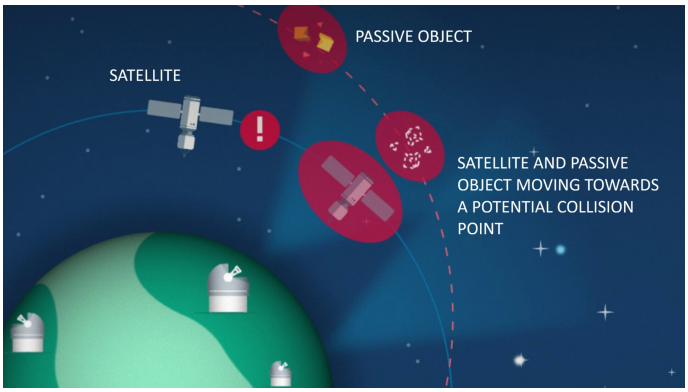

Figure 1.6: Dodging a passive object to keep the satellite safe (ESA Video)

short durations. However, modern mega-constellations dominantly utilize continuous, low-thrust electric propulsion [11, 31, 33]. While highly mass-efficient, low-thrust maneuvers present significant operational challenges: they require extended firing durations spanning several orbital revolutions, and their optimization inherently demands solving highly complex, continuous boundary value problems.

The significance of this work lies in addressing the mathematical and operational challenges of low-thrust CAMs. Previous research has demonstrated that while unconstrained energy-optimal continuous profiles can be analytically derived [16], finding the true fuel-optimal bang-bang sequence typically requires heavy numerical solvers or artificial directional thrust restrictions. By exploiting the empirical similarity in the costate vector profiles between energy- and fuel-optimal solutions, this thesis proposes a novel Differential Correction methodology implemented directly on the engine burning times. This approach combines the computational speed and 3D geometric fidelity of analytical energy-optimal formulations with the unconstrained mass-efficiency of a near-optimal bang-bang control law, offering a robust solution for autonomous onboard operations.

1.2. Research Status at Home and Abroad

The design of efficient CAMs is a critical requirement for spacecraft operations. Historically, astrodynamics research heavily focused on impulsive maneuvers modeled via analytical and numerical optimization [6, 8]. However, the widespread adoption of electric propulsion has fundamentally shifted the paradigm. Optimizing low-thrust CAMs necessitates the resolution of complex Two-Point Boundary Value Problems (TPBVP), prompting the development of diverse analytical, semi-analytical, and numerical strategies by researchers globally.

1.2.1. Analytical and Semi-analytical Approaches

To circumvent the heavy computational burden associated with the numerical integration of complex orbital dynamics, several authors have proposed analytical formulations by strategically simplifying the spacecraft dynamics or the optimal control problem itself. As an early baseline for impulsive maneuvers, Bombardelli and Hernando-Ayuso [8] established a fully analytical formulation that minimized the required Δv by reformulating the optimization task as an eigenvalue problem coupled with a nonlinear algebraic equation.

Transitioning to low-thrust applications, De Vittori et al. [16] developed a fully analytical algorithm specifically tailored for energy-optimal CAMs in short-term encounters. By projecting three-dimensional conjunction dynamics onto the B-plane and minimizing the control effort integral, they successfully derived closed-form solutions that tightly constrain either the collision probability or the deterministic physical miss distance. This robust energy-optimal framework was subsequently expanded by the same authors to incorporate complex post-maneuver requirements, designing Point-to-Point (PTP) and Point-to-Orbit (PTO) return strategies via a Three-Point Boundary Value Problem formulation [17]. Furthermore, this approach was elegantly adapted using Equinoctial Orbital Elements (EOE)—which explicitly account for $J_{2,2}$ geopotential perturbations—to analytically combine CAMs with long-term station-keeping operations for Geostationary (GEO) satellites [18, 20, 21].

Other researchers have sought analytical tractability by simplifying the relative dynamics equations governing the encounter. Ströbel and Stoll [55] utilized Hill-Clohessy-Wiltshire (HCW) equations to model the relative motion between the primary and secondary objects. By truncating the analytical series expansion of the collision probability and assuming low-thrust maneuvers primarily affect the along-track direction, they established a direct mathematical relationship to calculate the required burn and coasting durations.

In the pursuit of fuel-optimal solutions, a common analytical shortcut is to artificially constrain the thrust direction. Hernando-Ayuso and Bombardelli [27] derived an optimum continuous low-thrust maneuver for circular orbits by asserting a purely tangential thrust vector, thereby simplifying the task into a more manageable two-dimensional boundary value problem. Gonzalo et al. [25] enhanced this tangential model by explicitly retaining first-order terms in the thrust acceleration using averaging techniques, which accurately captures phasing changes and reduces model error without significantly increasing the computational burden. Conversely, for extremely short-warning encounters (less than half an orbital period), Reiter and Spencer [51] demonstrated via finite-burn analysis and multiple linear regression that a purely radial thrust direction becomes the optimal

semi-analytical solution.

Benefits and Challenges: As emphasized by recent comparative studies [43], the primary benefit of analytical and semi-analytical methods is their extraordinary computational speed. These methods are highly efficient, consistently running at least an order of magnitude faster than full numerical integrations, making them highly attractive for onboard autonomy and rapid ground-based screening. However, these benefits come with significant trade-offs. Semi-analytical low-thrust methods often exhibit higher required Δv because their foundational methodological constraints (e.g., forcing a binary choice between purely tangential or radial firing directions) prevent the optimizer from exploiting more fuel-efficient, unconstrained intermediate thrust vectors. Furthermore, fully analytical impulsive methods like [8] can sometimes underestimate the necessary control effort, leading to severe accuracy degradation and unacceptable safety constraint violations at the time of closest approach when directly applied to finite-burn engines.

1.2.2. Numerical Optimization Approaches

Solving the unconstrained Fuel-Optimal Control Problem (FOCP) typically mandates numerical methods due to the discontinuous, bang-bang nature of the optimal control law dictated by Pontryagin's Maximum Principle. In the broader context of low-thrust trajectory optimization, Wang and Grant [59] presented a foundational sequential convex programming (SCP) approach to efficiently solve the minimum-fuel orbit transfer problem. By applying a change of variables to decouple control parameters from system states, they relaxed inherently non-convex control bounds into convex second-order cone inequality constraints using a lossless convexification technique—mathematically proven to retain the exact optimal solution. By addressing non-linear orbital dynamics via successive small-disturbance linearizations, they successfully transformed the traditionally intractable optimal control problem into a sequence of rapidly converging second-order cone programming (SOCP) subproblems.

Building upon complex numerical integrations, Wang and Topputo [58] addressed fuel-optimal low-thrust transfers using indirect optimization, successfully navigating ill-conditioned State Transition Matrices and shadow eclipses via sophisticated energy-to-fuel continuation schemes. Dell'Elce et al. [22] bypassed complex shooting problems by backward integrating extremal trajectories to analytically map the cut and conjugate loci, thereby characterizing minimum warning times. Lee et al. [30] eliminated the need for initial guesses by embedding a fractional penalty function into the performance index, applying canonical transformation theory to produce a near-optimal feedback control law.

Recent state-of-the-art advancements in CAM design have increasingly leveraged these principles of Differential Algebra (DA) and convex optimization to manage the severe non-linearities of the FOCP. Armellin [6] modeled maneuvers as a sequence of multiple impulses, successfully solving them via SOCP and lossless convexification of the thrust magnitude. Pavanello et al. [41] extended convex optimization to handle multiple simultaneous short-term encounters by splitting spatial uncertainty using Gaussian Mixture Models (GMM) and applying SCP with dynamic trust regions. Additionally, Pavanello et al. [44] proposed a highly efficient recursive polynomial method that utilizes DA to approximate the Probability of Collision (PoC) as a high-order Taylor polynomial. This transforms the complex maneuver design into a sequence of rapid polynomial evaluations, bypassing heavy numerical solvers while accommodating non-linear safety constraints.

Pushing the boundaries of convexification further for collision avoidance, Pavanello et al. [45] recently introduced an advanced SCP technique specifically tailored for spacecraft engaged in continuous low-thrust propulsion. To navigate the highly nonlinear optimal control landscape, they utilized differential algebra to convexify the orbital dynamics and effectively approximated the probability of collision limit as a linearized squared Mahalanobis distance constraint. Addressing the growing challenge of dense debris environments, Zhou and Li [63] similarly proposed a multi-target collision avoidance strategy utilizing SCP to mitigate short-term encounters with several space objects simultaneously. Their methodology translates the complex continuous-thrust control problem into a constrained impulsive thrust planning framework, convexifying the non-linear collision probability limits to efficiently generate low-fuel maneuvers that ensure safety against multiple targets.

In scenarios involving non-cooperative targets or pursuit-evasion dynamics, intelligent search algorithms and machine learning have proven highly effective. Cao et al. [10] addressed turn-based, limited-area impulsive collision avoidance by proposing the Selection Probability Learning Adaptive Search-depth Search Tree (SPL-ASST) algorithm. By enhancing traditional Monte Carlo Tree Search (MCTS) with adaptive horizons and neural network-driven action selection, their model explicitly accounts for orbit determination delays, enabling rapid on-board generation of sequential multi-impulse evasion strategies. Furthermore, the management of severe risk uncertainty in multi-debris encounters has driven the adoption of advanced metaheuristic solvers. Lu et al. [35] applied Dempster-Shafer evidence theory to fuse heterogeneous safety metrics—such as multiple collision probabilities, Time of Closest Approach (TCA), and Distance of Closest Approach (DCA)—into a unified risk ranking. Leveraging this fused risk model, they deployed an Improved Memetic Algorithm (IMA) featuring rule-based initialization and

elite-guided local search to optimally balance spacecraft fuel consumption against strict multi-debris safety thresholds.

Seeking a computational middle ground between pure analytical and heavy numerical methods, De Vittori et al. [19] proposed a semi-analytical FOCP approach that extracts candidate bang-bang firing windows directly from continuous energy-optimal solutions. However, to maintain evaluation efficiency within a Non-Linear Programming (NLP) framework, this method strictly restricts the thrust execution to the tangential direction.

Benefits, Challenges, and the Literature Gap: The core benefit of numerical, optimization-based, and modern metaheuristic methods is their exceptional accuracy and unconstrained fuel efficiency. Comprehensive comparative analyses, such as the *CAMmary* review by Pavanello et al. [43] and sensitivity studies by Colombo et al. [12], demonstrate that methods leveraging DA, convex optimization [41, 44, 45, 63], and advanced evolutionary algorithms [35] are profoundly reliable. For instance, recursive polynomial methods and advanced SCP solvers can achieve spatial boundary violations of mere millimeters when targeting a specific Miss Distance, while consistently remaining within strict relative error limits for probabilistic PoC targets.

However, this immense precision comes at a steep computational cost. These algorithms require highly complex mathematical frameworks, heavy numerical integrations, or extensive neural network training pipelines [10], posing significant challenges for rapid, autonomous onboard implementation on satellites with limited processing capabilities. Furthermore, while the semi-analytical NLP method [19] is computationally lighter, the *CAMmary* analysis explicitly notes that it incurs the highest Δv penalty among low-thrust methods due to its strict tangential thrust constraint.

This highlights a distinct and persistent methodological gap within astrodynamics literature. While previous research [16] provided empirical proof that the time evolution of the velocity costate vector in a true fuel-optimal solution exhibits a striking similarity to the unconstrained 3D energy-optimal solution, no existing methodology successfully leverages this geometric duality to extract a near fuel-optimal bang-bang sequence without resorting to heavy numerical solvers, heuristic search trees, or artificial directional constraints. Specifically, the literature lacks a low-thrust solution that implements a differential correction algorithm directly on the burning times (t_{on} and t_{off}) starting from an accurate energy-optimal costate evolution, while imposing absolutely no constraints on the thrust direction. Bridging this gap to provide a fast, 3D unconstrained, near fuel-optimal solution natively suited for onboard execution is the objective of this thesis.

1.3. Main Content and Structure of the Thesis

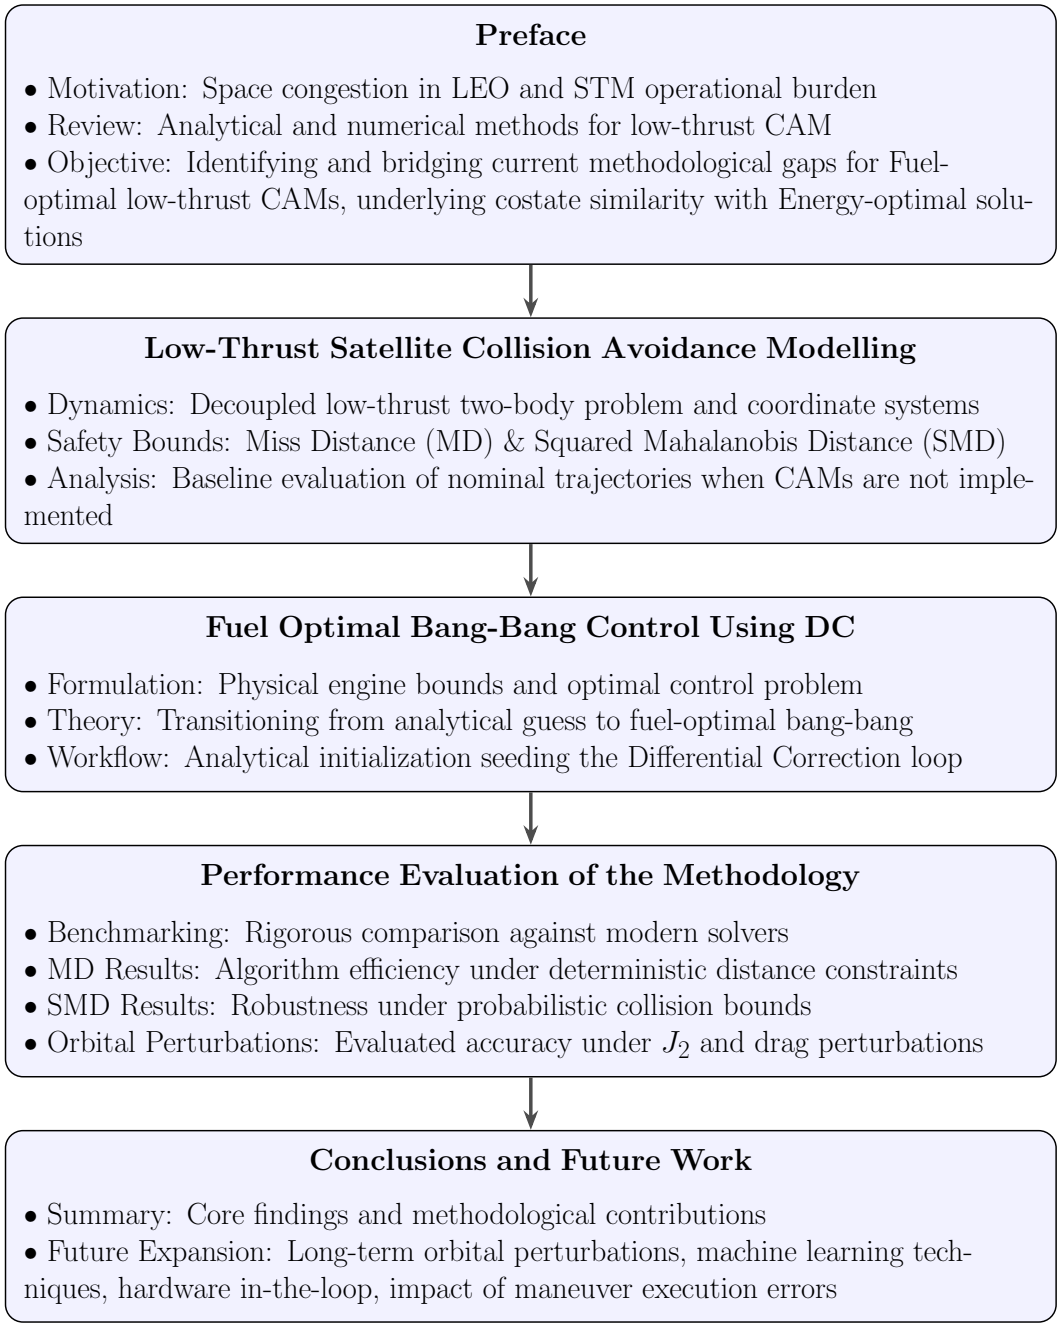

Figure 1.7: Diagram of the thesis content

This thesis focuses on developing a fast, robust, and autonomous methodology for designing near fuel-optimal low-thrust Collision Avoidance Maneuvers (CAMs) for short-term orbital encounters. By bridging the gap between highly efficient analytical methods and

strictly optimal numerical solvers, this research proposes a novel Differential Correction approach. The thesis is logically organized as follows, with block diagram available at Figure 1.7:

Chapter 1: Preface. This chapter introduces the motivation behind the research. Section 1.1 outlines the growing threat of space congestion in Low Earth Orbit (LEO) and the operational burden it places on Space Traffic Management. Section 1.2 critically reviews the current state-of-the-art analytical and numerical methods for low-thrust CAM optimization, explicitly identifying the methodological gaps that this thesis aims to resolve.

Chapter 2: Low-Thrust Satellite Collision Avoidance Modelling. This chapter establishes the fundamental physical, geometric, and mathematical models required for the study. Section 2.2 introduces the coordinate systems and the decoupled low-thrust two-body problem used to model the satellite dynamics. Section 2.3 defines the mathematical evaluation of collision conditions, formulating safety constraints based on both the Miss Distance and the Probability of Collision (scaled via the Squared Mahalanobis Distance). Finally, Section 2.4 provides a baseline statistical analysis of the dataset assuming nominal trajectories without avoidance maneuvers, explicitly highlighting the severity of the operational environment.

Chapter 3: Fuel Optimal Bang-Bang Control Using Differential Correction. This chapter presents the core methodological contribution of the thesis. Section 3.2 establishes the physical bounds of the propulsion system and formulates the optimal control problem. Section 3.3 details the theoretical transition from an analytically solvable energy-optimal control guess to a strict, fuel-optimal bang-bang profile. Building upon this foundation, Section 3.4 outlines the complete operational algorithm, demonstrating how the analytical initialization elegantly seeds the numerical Differential Correction loop to iteratively adjust the engine's ignition and cutoff times.

Chapter 4: Performance Evaluation of the Low-Thrust Fuel Optimal Differential Correction. This chapter evaluates the operational efficiency, precision, and fuel-optimality of the proposed methodology by rigorously benchmarking it against modern state-of-the-art solvers. Section 4.2 analyzes the algorithm's performance when subjected to deterministic physical miss distance constraints. Section 4.3 subsequently demonstrates its capability to robustly handle complex, probabilistic squared Mahalanobis distance bounds. Finally, Section 4.4 validates the operational robustness of the derived control profiles by assessing their constraint adherence under a high-fidelity dynamic environment incorporating atmospheric drag and J_2 Earth oblateness.

Chapter 5: Conclusions and Future Work. The final chapter summarizes the pri-

mary findings and contributions of this research and outlines logical avenues for future expansions, such as integrating high-fidelity orbital perturbations directly into the dynamics, employing machine learning techniques for algorithm initialization, testing Hardware-in-the-loop solutions and investigating the impact of maneuver execution errors.

2 | Low-Thrust Satellite Collision Avoidance Modelling

2.1. Introduction

This chapter establishes the fundamental physical, geometric, and mathematical models necessary for analyzing short-term orbital conjunctions and designing low-thrust collision avoidance maneuvers (CAMs). It begins by defining the primary coordinate systems and the dynamic models governing both the active primary satellite and the passive secondary object under a low-thrust two-body framework. Following the dynamic definitions, the chapter details the rigorous evaluation of collision conditions, elaborating on critical safety metrics such as the Miss Distance and the Probability of Collision, and how these are extracted from standardized Conjunction Data Messages (CDMs). To contextualize the operational necessity of these maneuvers, a baseline simulation of nominal encounters—without any avoidance actions—is presented, highlighting the high-density threat zones in modern Space Traffic Management. Finally, to bridge the gap between nonlinear orbital dynamics and computationally efficient targeting algorithms, the concept of the State Transition Matrix (STM) and perturbed solutions is derived. This sequence of theoretical foundations forms the mathematical bedrock for the optimal control strategies and differential correction methods explored in subsequent chapters.

2.2. Dynamics of the Satellites: Low-Thrust Two-Body Problem

2.2.1. Coordinate Systems

To rigorously describe the motion of a spacecraft and accurately formulate the collision avoidance problem, it is essential to define the appropriate reference frames. The dynamics and kinematics of the resident space objects are primarily analyzed using two fundamental coordinate systems: the global Earth-Centered Inertial frame and the local

Radial-Transversal-Normal frame.

The Earth-Centered Inertial (ECI) frame serves as the fundamental inertial reference for plotting trajectories and propagating orbital dynamics. Its origin is fixed at the center of mass of the Earth. The principal X-axis points toward the vernal equinox, the Z-axis aligns with the Earth's rotational axis pointing towards the celestial North Pole, and the Y-axis completes the right-handed Cartesian triad on the equatorial plane. Because this frame does not rotate with the Earth, it is considered inertially fixed, making it the standard frame for integrating Newton's laws of motion. The state of a spacecraft in the ECI frame is defined by its position vector $\mathbf{r}$ and velocity vector $\mathbf{v}$.

While the ECI frame is ideal for absolute dynamic propagation, relative motion and mission constraints—such as thrust pointing directions or sensor fields of view—are much more naturally expressed in a local, moving coordinate system attached to the spacecraft. For this purpose, the Radial-Transversal-Normal (RTN) frame, sometimes referred to as the local vertical/local horizontal (LVLH) frame, is introduced. Figure 2.1 provides a clear visualization of the relationship between the global ECI frame and the local RTN frame.

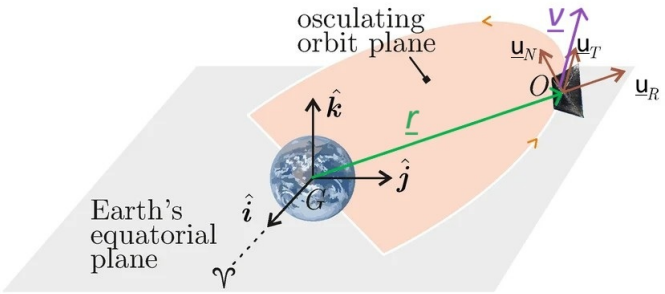

Figure 2.1: Visualization of ECI and RTN frames[37]

It is important to address a common geometrical misconception when constructing this local frame: while it is intuitively tempting to align the transversal axis directly with the spacecraft's velocity vector, the position and velocity vectors are only strictly perpendicular in perfectly circular orbits. In eccentric orbits, constructing a transformation matrix directly from the normalized position and normalized velocity vectors yields a non-orthogonal matrix, violating the fundamental mathematical properties of a rotation matrix.

To guarantee that the basis vectors form a mutually orthogonal triad, the RTN frame is rigorously defined using the orbital angular momentum. The three orthonormal unit vectors defining the RTN basis are constructed from the ECI position and velocity vectors

as follows:

$$\mathbf{u}_R = \frac{\mathbf{r}}{||\mathbf{r}||}$$

$$\mathbf{u}_N = \frac{\mathbf{r} \times \mathbf{v}}{||\mathbf{r} \times \mathbf{v}||}$$

$$\mathbf{u}_T = \mathbf{u}_N \times \mathbf{u}_R$$

Here, the Radial axis ($\mathbf{u}_R$) points outward along the position vector, the Normal axis ($\mathbf{u}_N$) points out of the orbital plane in the direction of the angular momentum, and the Transversal axis ($\mathbf{u}_T$) completes the right-handed system, pointing generally in the direction of motion.

To project vectors between these two coordinate systems, we utilize a 3×3 rotation matrix. Because the columns of a rotation matrix represent the unit vectors of the original frame expressed in the target frame, the transformation matrix from the RTN frame to the ECI frame is constructed simply by horizontally concatenating the RTN basis column vectors:

$$\mathbf{R}_{RTN}^{ECI} = \begin{bmatrix} \mathbf{u}_R & \mathbf{u}_T & \mathbf{u}_N \end{bmatrix}$$

If a vector is known in the RTN frame, denoted as $\mathbf{x}_{RTN}$, it can be transformed into the ECI frame via direct matrix multiplication:

$$\mathbf{x}_{ECI} = \mathbf{R}_{RTN}^{ECI} \mathbf{x}_{RTN}$$

Conversely, because $\mathbf{R}_{RTN}^{ECI}$ is a strictly orthogonal rotation matrix, its inverse is mathematically guaranteed to be equal to its transpose. Therefore, the transformation from the global ECI frame into the local RTN frame is achieved using the transposed matrix:

$$\mathbf{R}_{ECI}^{RTN} = \left(\mathbf{R}_{RTN}^{ECI}\right)^T = \begin{bmatrix} \mathbf{u}_R^T \\ \mathbf{u}_T^T \\ \mathbf{u}_N^T \end{bmatrix}$$

$$\mathbf{x}_{RTN} = \mathbf{R}_{ECI}^{RTN} \mathbf{x}_{ECI}$$

This rigorous mathematical mapping allows operators to define control bounds and physical miss distance constraints in the intuitive local RTN frame, while seamlessly transforming those parameters back into the ECI frame for continuous integration of the spacecraft's optimal control dynamics.

2.2.2. B-plane Definition

In a typical conjunction scenario, the satellite capable of performing maneuvers is designated as the "primary" object (O_p), whereas the unresponsive target—such as space debris or an inactive satellite—is identified as the "secondary" object (O_s). The states of these bodies are represented by the vectors $\mathbf{x}_p = (\mathbf{r}_p, \mathbf{v}_p)$ and $\mathbf{x}_s = (\mathbf{r}_s, \mathbf{v}_s)$. To better understand the relative geometry at the Time of Closest Approach (TCA), the encounter frame and the B-plane are illustrated in Figure 2.2.

To study the encounter, a reference frame is established at TCA, centered on the secondary object. The orthonormal basis vectors for this reference frame are constructed using the objects' velocities:

$$\mathbf{u}_x = \frac{\mathbf{v}_p}{||\mathbf{v}_p||}, \quad \mathbf{u}_z = \frac{\mathbf{v}_p \times \mathbf{v}_s}{||\mathbf{v}_p \times \mathbf{v}_s||}, \quad \mathbf{u}_y = \mathbf{u}_z \times \mathbf{u}_x$$

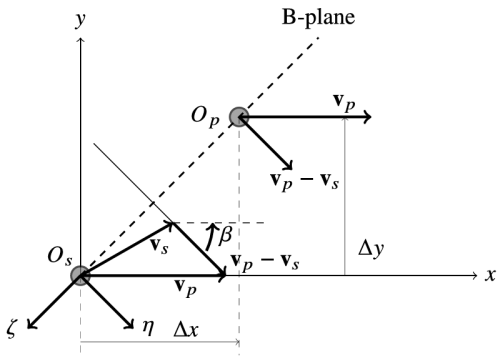

Figure 2.2: Encounter frame and B-plane: snapshot of $O_p - O_s$ encounter geometry[16]

An encounter is classified as "*short-term*" if its duration (t_c) is negligible compared to the primary object's orbital period (T_p), which mathematically occurs when their ratio is approximately 10^{-3} or less. Under this short-term hypothesis, the relative motion of the two bodies is treated as uniform and rectilinear. The dynamics of such collision avoidance scenarios are best analyzed using the B-plane, a specific plane perpendicular to the relative velocity of the two objects that contains both bodies at TCA. A 3D position vector within this coordinate system is denoted as $\mathbf{b}_{3D} = [\xi, \eta, \zeta]^T$. The basis vectors of the B-plane are derived from the relative velocity and the cross product of the primary and secondary velocities:

$$\mathbf{u}_\eta = \frac{\mathbf{v}_p - \mathbf{v}_s}{||\mathbf{v}_p - \mathbf{v}_s||}, \quad \mathbf{u}_\zeta = \frac{\mathbf{v}_s \times \mathbf{v}_p}{||\mathbf{v}_s \times \mathbf{v}_p||}, \quad \mathbf{u}_\xi = \mathbf{u}_\eta \times \mathbf{u}_\zeta$$

To map state variables from the standard ECI reference frame into the B-plane, the 3D rotation matrix $\mathbf{R}_{b,3D} = [\mathbf{u}_\xi, \mathbf{u}_\eta, \mathbf{u}_\zeta]^T$ is applied. If the problem is reduced to a purely two-dimensional projection on the B-plane axes, a truncated rotation matrix $\mathbf{R}_{b,2D} = [\mathbf{u}_\xi, \mathbf{u}_\zeta]^T$ is utilized instead.

2.2.3. Dynamics of the Primary Satellite

Given that the proposed collision avoidance methodology addresses short-term conjunctions with a maximum temporal horizon of a single orbit (detailed in Chapter 4), the cumulative effects of environmental perturbations like J_2 and atmospheric drag are considered negligible over this brief timeframe and are therefore excluded from the foundational dynamics of both primary and secondary satellites.

The primary satellite assumes an active role in the collision dynamics. Suppose we are analyzing the trajectory of this maneuverable spacecraft equipped with a low-thrust propulsion system, operating under the classical two-body dynamical model. The state of the primary spacecraft is initially fully described by a seven-dimensional state vector $\mathbf{x} = [\mathbf{r}^T, \mathbf{v}^T, m]^T$, which encompasses its three-dimensional position, velocity, and instantaneous mass. The equations of motion governing this system incorporate both the conservative gravitational acceleration and the non-conservative propulsive acceleration:

$$\begin{cases} \dot{\mathbf{r}} = \mathbf{v} \\ \dot{\mathbf{v}} = -\frac{\mu}{r^3}\mathbf{r} + u\frac{T_{max}}{m}\hat{\alpha} \\ \dot{m} = -u\frac{T_{max}}{I_{sp}g_0} \end{cases}$$

The propulsive capabilities are defined by the hardware parameters: the maximum available thrust T_{max}, the engine's specific impulse I_{sp}, and the standard gravitational acceleration g_0. The actual thrust applied at any given instant is dictated by two control variables: the thrust pointing unit vector $\hat{\alpha}$, which determines the thrust direction, and the dimensionless throttle factor $u \in [0, 1]$, which modulates the engine's thrust magnitude.

To streamline the optimal control formulation and facilitate numerical resolution, this study explicitly decouples the position and velocity dynamics from the mass dynamics. This mathematical simplification is highly justified for short-term collision avoidance maneuvers, which typically consume only a minuscule fraction of the onboard propellant. Because mass variations during such brief operational windows are functionally negligible, treating the spacecraft's mass as a constant ($m \approx m_0$) is a sound assumption that

conveniently yields a constant maximum acceleration capability. By applying this constant mass assumption, the state vector is reduced to six purely kinematic dimensions, $\mathbf{x} = [\mathbf{r}^T, \mathbf{v}^T]^T$, and the simplified equations of motion become:

$$\begin{cases} \dot{\mathbf{r}} = \mathbf{v} \\ \dot{\mathbf{v}} = -\frac{\mu}{r^3}\mathbf{r} + \mathbf{a}_c \end{cases}$$

where $\mathbf{a}_c = u a_{max} \hat{\alpha}$ is the commanded control acceleration vector, and $a_{max} = T_{max}/m_0$ represents the constant maximum acceleration evaluated at the initial maneuver mass.

Under this decoupled framework, the propellant mass expelled during the maneuver is evaluated strictly *a posteriori*, after the required trajectory optimization has concluded. Once the optimal control algorithm yields the continuous acceleration profile $\mathbf{a}_c(t)$, the total required velocity change (Δv) is obtained through direct temporal integration:

$$\Delta v = \int_{t_0}^{t_f} ||\mathbf{a}_c(t)|| \, dt \tag{2.1}$$

The corresponding consumed propellant mass, Δm, is then rigorously calculated using the classical Tsiolkovsky rocket equation:

$$\Delta m = m_0 \left[1 - \exp\left(-\frac{\Delta v}{I_{sp} g_0} \right) \right]$$

This decoupled approach drastically reduces the dimension and complexity of the subsequent Two-Point Boundary Value Problems (TPBVP) without sacrificing the physical validity or accuracy of the collision avoidance maneuver design.

2.2.4. Dynamics of the Secondary Satellite

The secondary satellite plays a passive role in the collision dynamics. It is assumed to be an uncooperative or non-maneuverable space object, such as a piece of space debris or a defunct payload. Therefore, its trajectory is governed purely by the classical two-body dynamical model without any propulsive acceleration. The equations of motion for the secondary satellite are described by the six-dimensional state vector $\mathbf{x}_s = [\mathbf{r}_s^T, \mathbf{v}_s^T]^T$:

$$\begin{cases} \dot{\mathbf{r}}_s = \mathbf{v}_s \\ \dot{\mathbf{v}}_s = -\frac{\mu}{r_s^3}\mathbf{r}_s \end{cases}$$

where $\mathbf{r}_s$ and $\mathbf{v}_s$ are the position and velocity vectors of the secondary satellite, respectively, and μ is the Earth's standard gravitational parameter.

2.3. Evaluation of Collision Conditions

2.3.1. Miss Distance

In the context of collision avoidance, the miss distance represents the closest physical proximity between the centers of mass of the primary and secondary objects during their encounter. Because the conjunction dynamics are projected onto the B-plane, the miss distance d is geometrically defined as the Euclidean norm, or magnitude, of the relative position vector on this plane. Let the two-dimensional relative position vector on the B-plane at the time of closest approach be denoted by $\mathbf{b}_f = [\xi_e, \zeta_e]^T$. The miss distance is mathematically expressed as:

$$d = \sqrt{\xi_e^2 + \zeta_e^2} = ||\mathbf{b}_f||$$

Alternatively, it can be calculated directly from the three-dimensional Cartesian relative position vectors of the primary ($\mathbf{r}_f$) and secondary ($\mathbf{r}_s$) objects. By applying the two-dimensional B-plane rotation matrix $\mathbf{R}_{b,2D}$ to project this difference onto the B-plane, the miss distance is equivalent to:

$$d = ||\mathbf{R}_{b,2D}(\mathbf{r}_f - \mathbf{r}_s)||$$

2.3.2. Probability of Collision and Squared Mahalanobis Distance

To rigorously evaluate the risk of a conjunction event, the probability of collision (PoC) between the maneuverable primary object (O_p) and the uncooperative secondary object (O_s) must be computed. Let the two space objects be modeled as spheres with radii s_1 and s_2, respectively. A collision occurs if the distance between their centers of mass becomes less than their combined cross-sectional radius, defined as $s_A = s_1 + s_2$. A visual representation of this combined cross-sectional radius is provided in Figure 2.3.

The calculation of the exact collision probability typically requires the integration of a three-dimensional relative position probability density function (pdf) over the duration of the encounter. However, under the short-term encounter hypothesis—where the rela-

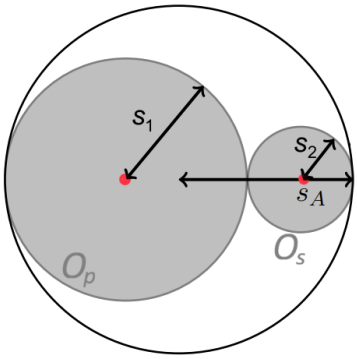

Figure 2.3: Visualization of combined cross-sectional radius s_A[39], where the centers of the two objects are depicted as red dots.

tive velocity is high and the time of the encounter is negligible compared to the orbital period—the relative motion can be approximated as uniform and rectilinear. This critical assumption allows the three-dimensional pdf to be integrated along the direction of relative motion, reducing the problem to a two-dimensional integral defined on the B-plane over a circular cross-section of radius s_A. Let the B-plane position vector of the primary relative to the secondary be denoted as $\mathbf{b} = [\xi, \zeta]^T$ (as previously seen in Figure 2.2). The relative position probability density function (pdf) is assumed to follow a bidimensional Gaussian distribution. The spatial uncertainty of the encounter is encapsulated by the combined 2D B-plane covariance matrix $\mathbf{C}$. Assuming the positional uncertainties of the two objects are statistically independent, this matrix is derived through a rigorous sequence of frame transformations and projections.

First, the 3×3 positional covariance matrices of the primary ($\mathbf{C}_p$) and secondary ($\mathbf{C}_s$) satellites are typically provided in their respective local Radial-Transverse-Normal (RTN) frames. Because these local frames have different orientations, the matrices must individually be rotated into a common inertial frame (the Earth-Centered Inertial, ECI, frame) using their specific rotation matrices, $\mathbf{R}_{RTN_p}^{ECI}$ and $\mathbf{R}_{RTN_s}^{ECI}$:

$$\mathbf{C}_{p,ECI} = \mathbf{R}_{RTN_p}^{ECI} \mathbf{C}_{p,RTN} \left(\mathbf{R}_{RTN_p}^{ECI}\right)^T \tag{2.2}$$

$$\mathbf{C}_{s,ECI} = \mathbf{R}_{RTN_s}^{ECI} \mathbf{C}_{s,RTN} \left(\mathbf{R}_{RTN_s}^{ECI}\right)^T \tag{2.3}$$

Once aligned in the common inertial space, the independent uncertainties are linearly summed to form the combined 3D covariance matrix:

$$\mathbf{C}_{comb,ECI} = \mathbf{C}_{p,ECI} + \mathbf{C}_{s,ECI}$$

Next, this combined spatial uncertainty is mapped onto the encounter geometry. It is rotated from the ECI frame into the 3D B-plane coordinate system utilizing the rotation matrix $\mathbf{R}_{b,3D}$. This matrix is constructed directly from the unit vectors defining the B-plane axes ($\mathbf{u}_\xi$, $\mathbf{u}_\eta$, and $\mathbf{u}_\zeta$):

$$\mathbf{R}_{b,3D} = [\mathbf{u}_\xi, \mathbf{u}_\eta, \mathbf{u}_\zeta]^T$$

$$\mathbf{C}_{comb,3D} = \mathbf{R}_{b,3D}\mathbf{C}_{comb,ECI}\mathbf{R}_{b,3D}^T$$

It is important to note that this combined covariance matrix inherently encapsulates the spatial uncertainties stemming from the Orbit Determination tracking processes, thereby natively embedding OD errors into the probabilistic safety constraints of the optimization problem.

Because the B-plane is strictly orthogonal to the relative velocity vector (the η-axis), the actual collision geometry is governed solely by the target plane axes (ξ and ζ). The final 2×2 covariance matrix $\mathbf{C}$ is obtained by applying a 2×3 Boolean projection matrix $\mathbf{P}$ to extract the components corresponding only to ξ and ζ:

$$\mathbf{C} = \mathbf{P}\mathbf{C}_{comb,3D}\mathbf{P}^T = \begin{bmatrix} \sigma_\xi^2 & \rho_{\xi\zeta}\sigma_\xi\sigma_\zeta \\ \rho_{\xi\zeta}\sigma_\xi\sigma_\zeta & \sigma_\zeta^2 \end{bmatrix}$$

where σ_ξ^2 and σ_ζ^2 are the spatial variances along the respective B-plane coordinates, and $\rho_{\xi\zeta}$ is the correlation coefficient between them. These fundamental elements define the shape and orientation of the combined covariance ellipse on the B-plane (as depicted in Figure 2.4).

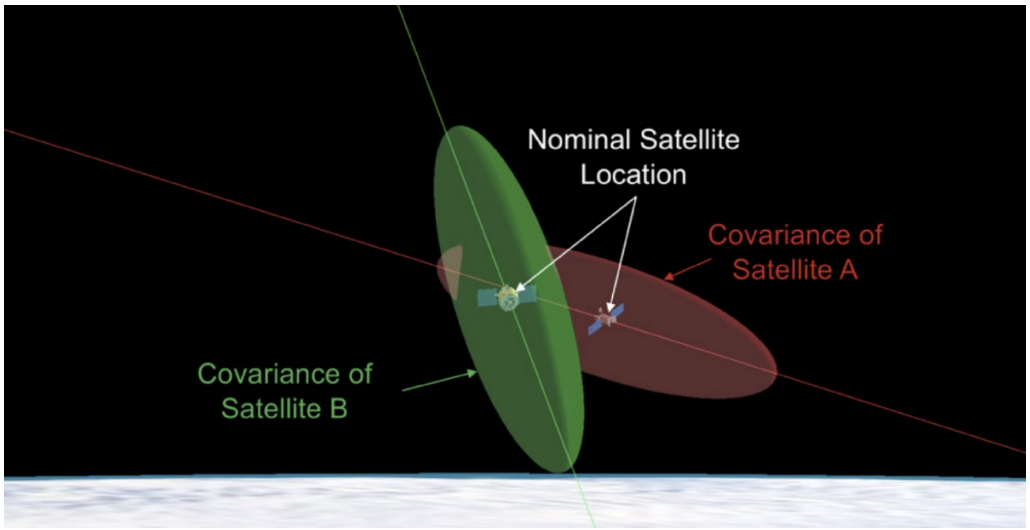

Figure 2.4: Visualization of covariances intersecting along orbital paths [2]

A standard approach to approximating this two-dimensional integral is Chan's method of equivalent cross-sectional areas. By approximating the integration domain as a circular cross-section of equal area, the probability of collision evaluates to a Rician integral, which can be computed using a convergent series expansion:

$$PoC(u,v) = e^{-v/2} \sum_{m=0}^{\infty} \frac{v^m}{2^m m!} \left[1 - e^{-u/2} \sum_{k=0}^{m} \frac{u^k}{2^k k!} \right]$$

The formulation relies on two fundamental parameters, u and v. The parameter u represents the ratio of the combined impact cross-sectional area to the area of the one-sigma covariance ellipse in the B-plane:

$$u = \frac{s_A^2}{\sigma_\xi \sigma_\zeta \sqrt{1 - \rho_{\xi\zeta}^2}}$$

The parameter v corresponds to the **squared Mahalanobis distance** (d_M^2) between the two objects at the time of closest approach. It serves as a non-dimensional measure of the miss distance, scaled by the encounter's covariance:

$$v = d_M^2 = \mathbf{b}_f^T \mathbf{C}^{-1} \mathbf{b}_f = \frac{(\xi/\sigma_\xi)^2 + (\zeta/\sigma_\zeta)^2 - 2\rho_{\xi\zeta}(\xi\zeta)/(\sigma_\xi\sigma_\zeta)}{1 - \rho_{\xi\zeta}^2}$$

Alternatively, a highly accurate evaluation can be obtained using the García-Pelayo and Hernando-Ayuso analytical formula, which provides a first-order expansion of the integral:

$$PoC \simeq \frac{s_A^2}{2\sigma_x\sigma_y} \left[1 + \frac{1}{8} \left(\frac{x^2/\sigma_x^2 - 1}{\sigma_x^2} + \frac{y^2/\sigma_y^2 - 1}{\sigma_y^2} \right) \right] \exp\left(-\frac{1}{2} \left(\frac{x^2}{\sigma_x^2} + \frac{y^2}{\sigma_y^2} \right) \right)$$

where x and y are the B-plane position vector components projected along the eigenaxes of the covariance ellipse, and σ_x and σ_y are the respective semi-principal axes. The relationship between the conjunction encounter and the resulting combined covariance in the encounter plane is further detailed in Figure 2.5.

From an optimization standpoint, it is crucial to note that the collision probability calculated via Chan's formulation decreases monotonically with respect to the squared Mahalanobis distance v. Because of this monotonic relationship, the squared Mahalanobis distance can be directly maximized as a surrogate cost function during the optimal control design of the collision avoidance maneuver, drastically simplifying the numerical resolution of the problem.

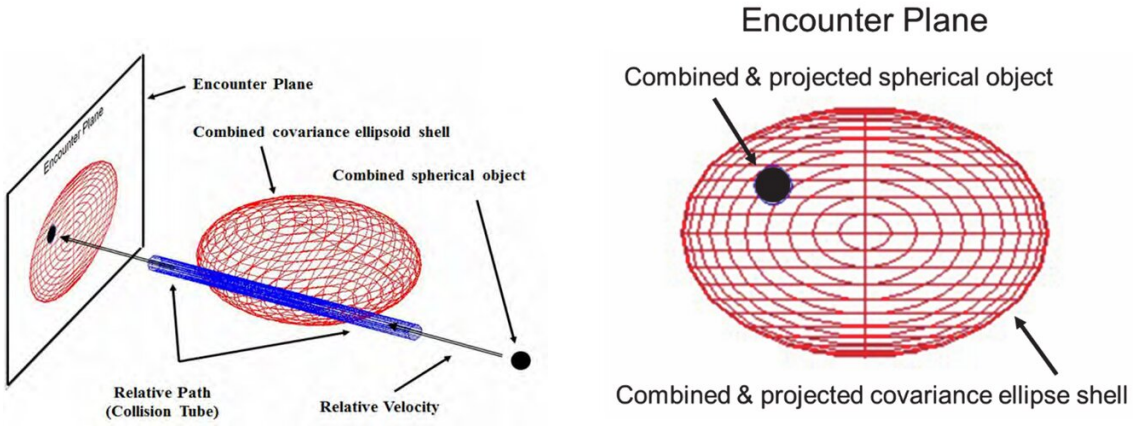

(a) Conjunction encounter visualization. (b) Combined covariance in encounter plane.

Figure 2.5: Conjunction encounter and combined covariance[5]

2.3.3. Conjunction Data Message

A conjunction in space operations refers to a predicted close approach between two resident space objects. When trajectories are propagated forward in time, a conjunction event is flagged if a secondary object, such as space debris or a non-operational satellite, enters a predefined safety screening volume surrounding a maneuverable primary satellite. The exact moment when the physical distance between these two bodies reaches its absolute minimum is known as the Time of Closest Approach (TCA). At TCA, the state vectors of the primary and secondary objects, denoted as $\mathbf{x}_p = (\mathbf{r}_p, \mathbf{v}_p)$ and $\mathbf{x}_s = (\mathbf{r}_s, \mathbf{v}_s)$ respectively, become the fundamental parameters for assessing the severity of the encounter. The relative state vector at this critical moment is simply the difference between the primary and secondary states:

$$\Delta\mathbf{x} = \mathbf{x}_p - \mathbf{x}_s$$

When a high-risk conjunction is detected by space surveillance networks, mission operators receive a Conjunction Data Message (CDM). A CDM is a standardized data format used globally to transmit comprehensive information about a predicted close approach [1]. It provides the essential parameters required to evaluate the collision risk and, if necessary, design an avoidance maneuver. The message contains the precise state vectors of both objects at TCA, along with their respective physical characteristics such as mass and estimated cross-sectional areas [1]. Crucially, the CDM also includes the positional and temporal uncertainty of both bodies. This uncertainty is mathematically quantified by individual 6×6 covariance matrices, $\mathbf{C}_p$ and $\mathbf{C}_s$, which capture the statistical errors inherent in the orbit determination and propagation process.

2.4. Simulation Results: Nominal Encounters Without CAM

Before introducing the optimal control logic for collision avoidance, it is necessary to establish a baseline by evaluating the dataset under a "Do-Nothing" scenario. In this baseline, the primary satellite (O_p) follows its nominal ballistic trajectory without executing any evasive maneuvers. Analyzing this unperturbed state highlights the severity of the current Space Traffic Management (STM) environment and justifies the need for advanced probabilistic modeling.

2.4.1. Construction of Avoidance Constraints

To automate the avoidance process within the optimal control framework, mathematical safety boundaries must be defined on the B-plane. The method of modeling these constraints fundamentally dictates the required control effort. Figure 2.6 provides a three-dimensional representation of the risk volume.

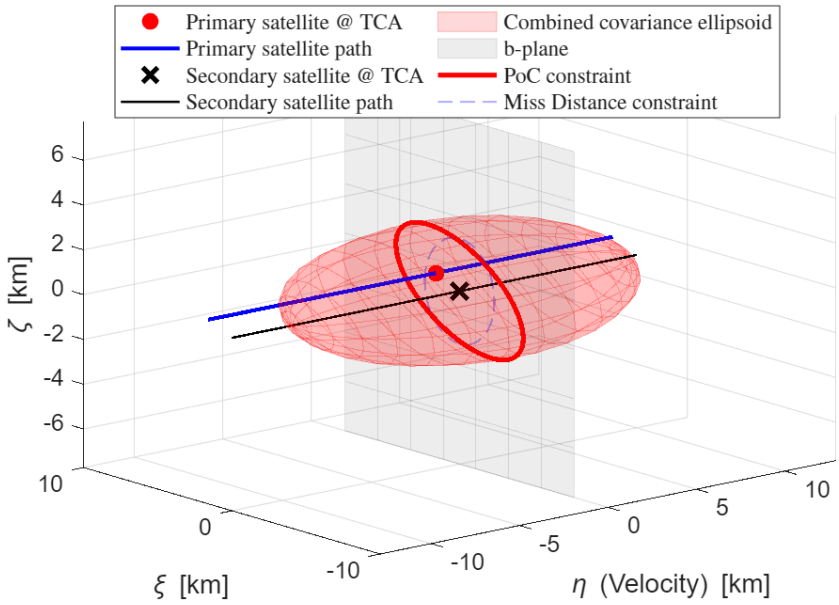

Figure 2.6: 3D Combined Covariance Ellipsoid

As illustrated in Figure 2.6, the true physical risk is represented by a 3D combined covariance ellipsoid. The relative velocity vector dictates the orientation of the B-plane, which acts as a 2D cross-section orthogonal to the incoming trajectory at TCA. Figure 2.7 compares how geometric and probabilistic constraints map onto this 2D plane.

Figure 2.7a demonstrates the construction of the miss distance constraint on the B-plane.

This geometric boundary generates a perfectly circular safety threshold, treating spatial uncertainty as uniform in all directions. While mathematically simple to implement in optimal control problems (as it relies purely on the Euclidean norm $||\mathbf{b}_f||$), it frequently overestimates risk in the cross-track direction while underestimating it along the velocity axis.

In contrast, Figure 2.7b shows the construction of the PoC constraint using the squared Mahalanobis distance (d_M^2). By projecting the 3D combined covariance matrix onto the 2D B-plane, the resulting constraint is a highly deformed ellipse.

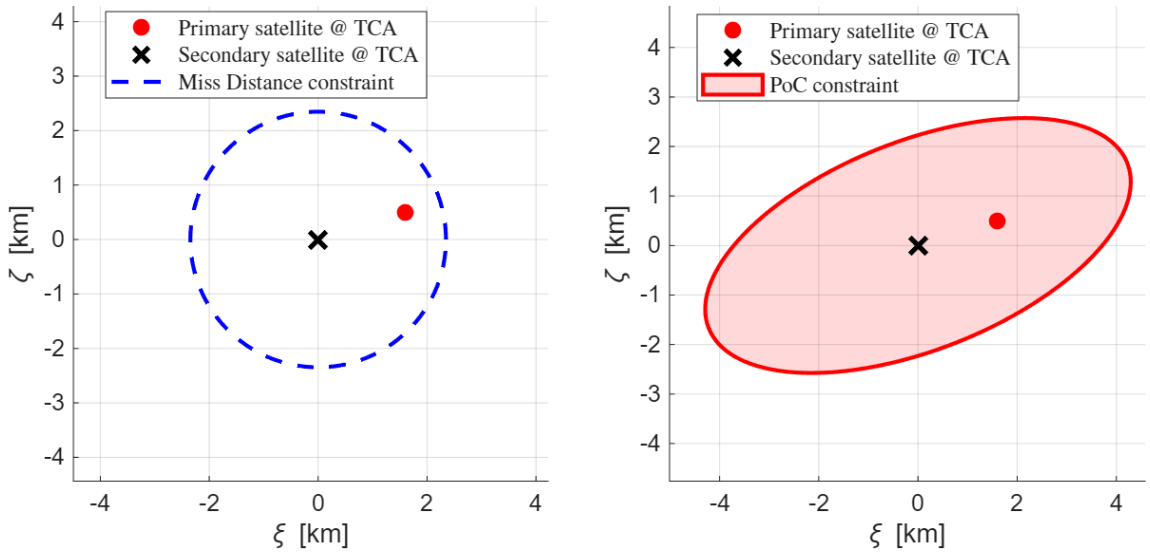

(a) Miss Distance Constraint (d). (b) Probabilistic Constraint (PoC).

Figure 2.7: Comparison of geometric and probabilistic collision avoidance constraints on the B-plane.

This footprint serves as the mathematically rigorous "shadow" of the 3D uncertainty volume. Resolving the collision avoidance problem therefore requires pushing the primary satellite's nominal B-plane crossing coordinate safely outside of this specific, covariance-scaled elliptical boundary.

2.4.2. Encounter Distribution and Heatmap Generation

To rigorously validate the proposed modelling approach, a comprehensive set of realistic test cases is utilized. This dataset is sourced from the official Collision Avoidance Challenge organized by the European Space Agency (ESA) and is publicly available for download[1]. To visualize the spatial distribution of the collision risks, the 3D relative state

[1] https://github.com/arma1978/conjunction

vectors of all encounters at the Time of Closest Approach (TCA) are projected onto a normalized 2D B-plane. The secondary objects (O_s) are standardized at the origin of this plane. The heatmaps are generated using 2D histograms to capture the encounter density, highlighting the highest concentration zones. Figure 2.8 presents the distribution of unperturbed encounters evaluated against two different safety metrics introduced earlier in the chapter: the miss distance (d) and the Probability of Collision (PoC).

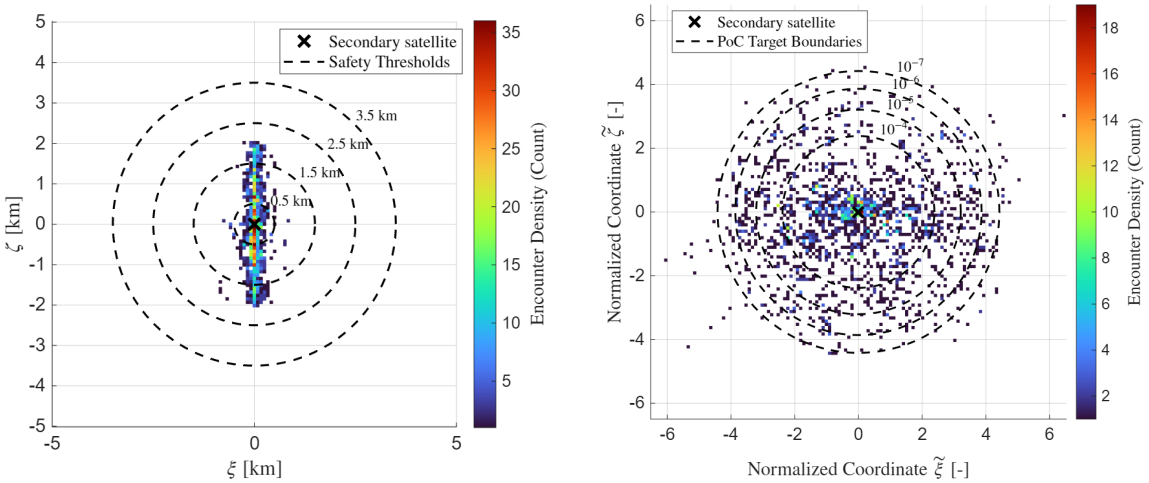

(a) Miss Distance (d) Heatmap.

(b) Probability of Collision (PoC) Heatmap.

Figure 2.8: Normalized B-plane heatmaps of all dataset samples under nominal conditions.

As shown in Figure 2.8a, when evaluated strictly in terms of Euclidean geometry on the B-plane, the dataset is highly constrained. Almost all unperturbed samples yield a miss distance strictly below 2 km, creating a dense, highly concentrated threat zone immediately surrounding the secondary satellite. Conversely, Figure 2.8b visualizes the same dataset transformed into the Mahalanobis space, where physical distances are scaled by the encounter-specific combined covariance matrices ($\mathbf{C}$). In this probabilistic domain, the distribution is significantly more evenly spread. While many encounters still violate the lenient 10^{-4} threshold, the data shows a broad dispersion reaching into the extreme high-risk 10^{-7} boundaries for several samples. This proves that an encounter with an ostensibly "safe" physical miss distance can still yield an unacceptably high PoC if the tracking uncertainty (covariance) is large or highly elongated along the B-plane eigenaxes.

The dense concentration of primary satellites encroaching upon the origin underscores a critical operational dilemma. Because these high-concentration zones lie dangerously close to the secondary satellite, the decision to execute a collision avoidance maneuver relies entirely on the owner/operator's internal risk assessment and chosen safety threshold. If an operator employs a simplistic spherical model (relying solely on a fixed target

miss distance, $\bar{d}$) or sets an excessively lenient PoC threshold to save propellant, they may choose to bypass a maneuver. However, such irresponsible operational behavior in densely populated orbital regimes (such as LEO) carries disastrous potential. A single catastrophic collision resulting from a neglected high-risk conjunction generates thousands of untrackable lethal debris fragments. This localized fragmentation can quickly cascade—triggering the Kessler Syndrome—ultimately rendering the orbital plane hazardous for all other satellites and constellations.

2.5. Summary of This Chapter

This chapter establishes the fundamental theoretical framework and mathematical models required to analyze short-term orbital conjunctions and design low-thrust collision avoidance maneuvers.

1) First, the geometric and probabilistic environments of a conjunction event are defined. By leveraging Conjunction Data Messages (CDMs) and transitioning from the global Earth-Centered Inertial (ECI) frame to the local Radial-Transversal-Normal (RTN) frame and the encounter-specific B-plane, the formulation of critical safety constraints is formalized. These constraints encompass both the deterministic physical miss distance and the statistical Probability of Collision (P_c), which is mathematically represented by the covariance-scaled squared Mahalanobis distance (d_M^2).

2) Subsequently, the dynamics of the resident space objects are modeled using a decoupled low-thrust two-body problem. This framework treats the primary satellite's mass variation as negligible over the short maneuver window, significantly simplifying the optimization process.

3) The chapter then presents an analysis of the ESA dataset under a nominal, no-maneuver baseline scenario. The generated B-plane heatmaps visualize the high-density threat zones surrounding the secondary objects and underscore the critical operational differences between evaluating purely geometric constraints versus probabilistic constraints. This baseline evaluation explicitly highlights the severity of the current space traffic environment and solidifies the operational necessity for robust, automated, and highly mass-efficient collision avoidance algorithms.

The following section introduces the state transition matrix, a critical tool required to map the effects of maneuvers across these complex orbital geometries.

3 | Fuel Optimal Bang-Bang Control Using Differential Correction

3.1. Introduction

This chapter establishes the mathematical models, dynamical frameworks, and optimal control strategies utilized to design low-thrust collision avoidance maneuvers (CAMs). While the preceding chapters lay the groundwork by defining the collision geometry, the decoupled two-body dynamics, and the linear mapping properties of perturbed trajectories, this chapter focuses on the practical synthesis of the maneuver itself. The subsequent sections detail the transformation of the spacecraft's physical operational limits into rigorous boundary value problems, demonstrating how optimal continuous and discontinuous thrust profiles are computationally derived to satisfy stringent mission safety constraints.

The fundamental challenge in designing a low-thrust CAM lies in the continuous nature of the propulsion system. Unlike traditional chemical engines that provide near-instantaneous velocity changes (Δv), electric propulsion requires extended firing arcs that span significant fractions of the orbital period. To maximize the operational lifespan of the satellite, it is imperative to formulate this as a Fuel-Optimal Control Problem (FOCP). This chapter details the evolution from an analytically tractable Energy-Optimal Control Problem (EOCP) to a highly robust, numerically stable Differential Correction (DC) methodology that explicitly targets true fuel optimality via a discontinuous bang-bang control law.

3.2. Problem Formulation

The primary objective of a fuel-optimal collision avoidance maneuver is to guarantee that the primary satellite achieves a predefined safety target at the Time of Closest Approach (TCA) while consuming the absolute minimum amount of onboard propellant. The safety

target is formulated either as a deterministic physical MD or a probabilistic boundary scaled by the encounter's tracking uncertainty (the SMD).

In real-world spaceflight operations, a fuel-optimal solution is strictly preferred over an energy-optimal one due to practical hardware limitations. An energy-optimal control profile dictates a continuous, varying thrust magnitude over the entire maneuver window. However, this often yields commanded thrust values that are excessively low, which causes electric propulsion engines to unintentionally shut off or operate outside their stable design parameters. Furthermore, continuously throttling an engine over a prolonged time interval is highly susceptible to accumulating execution and navigation errors.

Conversely, a fuel-optimal solution operates exclusively at the engine's physical extremes, demanding a strict "bang-bang" sequence. Operating the engine discretely—either firing at its maximum capability (100% throttle) or completely shut off (0% throttle)—is easily achievable, highly reliable, and vastly more precise for modern propulsion systems. Additionally, by confining the thrust to a shorter, concentrated time span, the bang-bang maneuver inherently minimizes the accumulation of trajectory execution errors compared to a continuous, orbit-long thrusting solution.

Because the spacecraft's engine operates at this fixed maximum thrust capacity (a_{max}), minimizing the total fuel consumption is mathematically equivalent to minimizing the total duration the engine is actively firing. According to Pontryagin's Maximum Principle (PMP), the optimal thrust profile for such a bounded control problem naturally corroborates this strict "bang-bang" sequence, with no intermediate states permitted. Consequently, the trajectory optimization problem is entirely reduced to identifying the exact optimal moments to ignite the engine (the ignition time, t_{on}) and to shut it down (the cutoff time, t_{off}).

Traditionally, resolving these switching times requires solving a highly non-linear Two-Point Boundary Value Problem (TPBVP) by integrating sensitive costate differential equations. This standard approach is notoriously susceptible to poor initial guesses and frequently suffers from convergence failures when attempting to model discontinuous pulse functions.

To circumvent this severe computational bottleneck, this chapter introduces a hybrid methodology. The strategy leverages the closed-form analytical solutions of the Energy-Optimal Control Problem to define the underlying geometry and thrust direction of the maneuver. These energy-optimal parameters are then used to seed a Differential Correction loop. Rather than integrating volatile costates, the DC method relies on the State Transition Matrix (STM)—a linear sensitivity matrix that evaluates how perturbations

at a specific epoch amplify or decay over time that maps variations in the switching times directly to the terminal constraint error at TCA, avoiding to solve a not trivial TPBVP. A visualization of the rationale behind the method is provided in Figure 3.1, and the method is described in the following section.

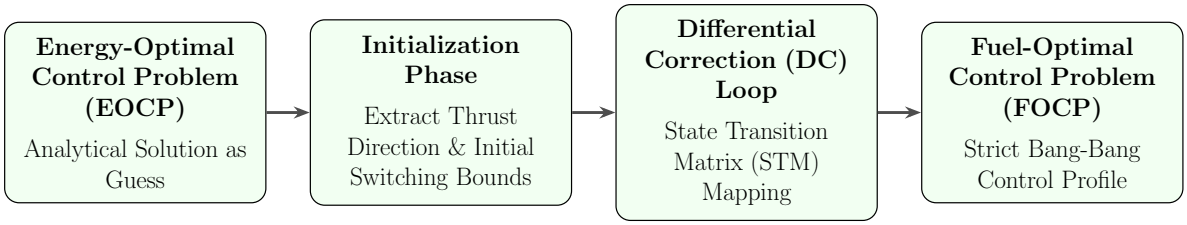

Figure 3.1: Diagram showing the rationale behind the method

3.3. Method Description

3.3.1. Energy Optimal Analytical Guess

The mathematical model and procedural framework detailed in this section are drawn directly from the work of De Vittori et al. [16]. This established approach serves as the theoretical baseline upon which the novel methodology proposed in this work is built.

In the context of low-thrust collision avoidance maneuver (CAM) design, formulating the optimization as an Energy-Optimal Control Problem (EOCP) provides an analytically tractable pathway to determine continuous thrust profiles. This section explores the controlled motion of the primary object within the Earth-Centered Inertial (ECI) coordinate system, governed by propelled Keplerian dynamics. The control variable in this formulation is the acceleration vector $\mathbf{a}_c = [a_{c,x}, a_{c,y}, a_{c,z}]^T$. To ensure numerical stability and facilitate optimal control derivation, the acceleration is parameterized as [16]:

$$\mathbf{a}_c = a_{max}\epsilon\alpha$$

where a_{max} represents a scaling parameter indicative of the maximum thrust acceleration available (assumed to be constant), ϵ is a dimensionless factor modulating the thrust magnitude, and $\alpha = \mathbf{a}_c/||\mathbf{a}_c||$ denotes the unit vector specifying the thrust direction.

The primary objective of the EOCP is to minimize the integral of the squared control effort over the maneuver duration. The cost functional is defined as:

$$J = \int_{t_0}^{t_f} \frac{1}{2} a_{max}\epsilon^2 dt$$

where t_0 and t_f are the initial maneuver time and the TCA, respectively.

Depending on the mission safety requirements, the terminal condition at TCA is constrained either by a predefined SMD or a specific physical MD. These two distinct boundary value problems are formulated and solved sequentially in the following subsections.

1. Square Mahalanobis Distance Constraint:

In this formulation, the primary object's relative position at conjunction, $\mathbf{r}_f = \mathbf{r}(t_f)$, is constrained to achieve a specific target SMD, $\bar{d}_M^2$. This statistically scales the MD by the combined positional uncertainty of both objects. The terminal boundary condition is rigorously stated as:

$$\Psi(\mathbf{x}(t_f), t_f) = d_M^2(\mathbf{r}_f) - \bar{d}_M^2 = 0$$

To solve this constrained optimization problem, the terminal condition is adjoined to the cost functional using a scalar Lagrange multiplier ν, yielding the augmented cost function [16]:

$$\bar{J} = \nu[d_M^2(\mathbf{r}_f) - \bar{d}_M^2] + \int_{t_0}^{t_f} \left\{ \frac{1}{2} a_{max} \epsilon^2 + \lambda^T(t)[\dot{\mathbf{x}}(t) - \mathbf{f}(\mathbf{x}, \mathbf{a}_c)] \right\} dt$$

where $\mathbf{f}$ encapsulates the nonlinear equations of motion, and $\lambda(t) = [\lambda_r^T, \lambda_v^T]^T$ represents the time-varying costate vector.

The corresponding Hamiltonian for this system is constructed as:

$$H = \frac{1}{2} a_{max} \epsilon^2 + \lambda_r^T \mathbf{v} + \lambda_v^T \left(-\frac{\mu}{r^3} \mathbf{r} + \epsilon a_{max} \alpha \right)$$

Applying Pontryagin's Minimum Principle allows us to extract the optimal control laws for the thrust magnitude and direction. The conditions that minimize the Hamiltonian dictate that the optimal thrust pointing vector α^* must directly oppose the velocity costate, and the optimal throttle fraction ϵ^* equates to the magnitude of the velocity costate [16]:

$$\begin{cases} \alpha^* = -\frac{\lambda_v}{||\lambda_v||} \\ \epsilon^* = ||\lambda_v|| \end{cases}$$

Consequently, the optimal unconstrained control acceleration evaluates to $\mathbf{a}_c^* = -a_{max} \lambda_v$. Incorporating this optimal control law, the Euler-Lagrange equations yield a coupled nonlinear Two-Point Boundary Value Problem (TPBVP) governing the state and costate

derivatives:

$$\begin{cases} \dot{\mathbf{r}} = \mathbf{v} \\ \dot{\mathbf{v}} = -\frac{\mu}{r^3}\mathbf{r} - a_{max}\lambda_v \\ \dot{\lambda}_r = \frac{\mu}{r^3}\lambda_v - \frac{3\mu}{r^5}(\mathbf{r}^T\lambda_v)\mathbf{r} \\ \dot{\lambda}_v = -\lambda_r \end{cases}$$

The boundary conditions mandate that the initial state is fixed ($\mathbf{r}(t_0) = \mathbf{r}_0$, $\mathbf{v}(t_0) = \mathbf{v}_0$), the final velocity is unconstrained leading to a null terminal velocity costate ($\lambda_v(t_f) = \mathbf{0}$), and the terminal position costate satisfies the transversality condition proportional to the gradient of the terminal constraint ($\lambda_r(t_f) = \nu\frac{\partial d_M^2}{\partial \mathbf{r}_f}$).

Because solving this TPBVP via standard numerical shooting methods is computationally expensive and highly sensitive to initial guesses, an analytical resolution is achieved by linearizing the dynamics about the nominal, uncontrolled (Keplerian) trajectory. This reduces the problem to an IVP mapped by the STM, $\mathbf{\Phi}(t, t_0)$.
The STM is obtained by integrating:

$$\dot{\mathbf{\Phi}}(t, t_0) = \mathbf{A}(t)\mathbf{\Phi}(t, t_0), \quad \mathbf{\Phi}(t_0, t_0) = \mathbf{I}_{12\times 12}$$

where $\mathbf{A}(t)$ is the Jacobian of the necessary conditions evaluated along the Keplerian motion. The STM maps variations from the initial state and costate to their terminal counterparts [16]:

$$\begin{bmatrix} \delta\mathbf{r}_f \\ \delta\mathbf{v}_f \\ \delta\lambda_{rf} \\ \delta\lambda_{vf} \end{bmatrix} = \begin{bmatrix} \mathbf{\Phi}_{11} & \mathbf{\Phi}_{12} & \mathbf{\Phi}_{13} & \mathbf{\Phi}_{14} \\ \mathbf{\Phi}_{21} & \mathbf{\Phi}_{22} & \mathbf{\Phi}_{23} & \mathbf{\Phi}_{24} \\ \mathbf{\Phi}_{31} & \mathbf{\Phi}_{32} & \mathbf{\Phi}_{33} & \mathbf{\Phi}_{34} \\ \mathbf{\Phi}_{41} & \mathbf{\Phi}_{42} & \mathbf{\Phi}_{43} & \mathbf{\Phi}_{44} \end{bmatrix} \begin{bmatrix} \delta\mathbf{r}_0 \\ \delta\mathbf{v}_0 \\ \delta\lambda_{r0} \\ \delta\lambda_{v0} \end{bmatrix}$$

Since the initial kinematic state is rigidly defined, $\delta\mathbf{r}_0 = \mathbf{0}$ and $\delta\mathbf{v}_0 = \mathbf{0}$. Along the nominal uncontrolled trajectory, the costates are inherently zero; thus, their variations are equivalent to the required optimal initial costates: $\delta\lambda_{r0} = \lambda_{r0}$ and $\delta\lambda_{v0} = \lambda_{v0}$.
At the final time, the velocity costate remains strictly zero ($\delta\lambda_{vf} = \mathbf{0}$). Utilizing the lower blocks of the STM, the initial velocity costate is expressed explicitly as a function of the initial position costate [16]:

$$\lambda_{v0} = -\mathbf{\Phi}_{44}^{-1}\mathbf{\Phi}_{43}\lambda_{r0}$$

Also, the initial position costate maps to the final position costate via a sensitivity matrix $\mathbf{B}$ [16]:

$$\lambda_{r0} = \mathbf{B}^{-1}\lambda_{rf} \quad \text{where} \quad \mathbf{B} = \mathbf{\Phi}_{33} - \mathbf{\Phi}_{34}\mathbf{\Phi}_{44}^{-1}\mathbf{\Phi}_{43}$$

The deviation in the terminal position achieved by the maneuver, $\delta\mathbf{r}_f = \mathbf{r}_f - \mathbf{r}_p$, is similarly

mapped to the terminal position costate using an auxiliary matrix $\mathbf{D}$ [16]:

$$\delta\mathbf{r}_f = \mathbf{D}\mathbf{B}^{-1}\lambda_{rf} \quad \text{where} \quad \mathbf{D} = \mathbf{\Phi}_{13} - \mathbf{\Phi}_{14}\mathbf{\Phi}_{44}^{-1}\mathbf{\Phi}_{43}$$

Substituting the terminal transversality condition $\lambda_{rf} = 2\nu\mathbf{R}_{b,2D}^T\mathbf{C}^{-1}\mathbf{R}_{b,2D}(\mathbf{r}_f - \mathbf{r}_s)$ into the state deviation equation bridges the 3D ECI dynamics to the 2D B-plane constraints. Pre-multiplying by the 2D B-plane rotation matrix $\mathbf{R}_{b,2D}$ and defining the B-plane vectors $\mathbf{b}_f = \mathbf{R}_{b,2D}(\mathbf{r}_f - \mathbf{r}_s)$ and $\mathbf{b}_p = \mathbf{R}_{b,2D}(\mathbf{r}_p - \mathbf{r}_s)$, the system is elegantly reduced to:

$$\begin{cases} \mathbf{b}_f - \mathbf{b}_p = 2\nu\mathbf{R}_{b,2D}\mathbf{D}\mathbf{B}^{-1}\mathbf{R}_{b,2D}^T\mathbf{C}^{-1}\mathbf{b}_f \\ \mathbf{b}_f^T\mathbf{C}^{-1}\mathbf{b}_f = \bar{d}_M^2 \end{cases}$$

To further condense this, we define a composite characteristic matrix[16]:

$$\mathbf{E} = 2\mathbf{R}_{b,2D}\mathbf{D}\mathbf{B}^{-1}\mathbf{R}_{b,2D}^T\mathbf{C}^{-1}$$

The final B-plane position is isolated as:

$$\mathbf{b}_f = (\mathbf{I}_{2\times2} - \nu\mathbf{E})^{-1}\mathbf{b}_p$$

Inserting this mapped position into the Mahalanobis boundary condition produces a deterministic scalar equation. By introducing the adjugate-like expression $\tilde{\mathbf{E}} = \det(\mathbf{E}) \cdot \mathbf{E}^{-1}$ [16], the problem collapses into a closed-form fourth-order polynomial in the unknown Lagrange multiplier ν [16]:

$$\nu^2(\tilde{\mathbf{E}}\mathbf{b}_p)^T\mathbf{C}^{-1}(\tilde{\mathbf{E}}\mathbf{b}_p) - \nu\left[\mathbf{b}_p^T\mathbf{C}^{-1}\tilde{\mathbf{E}}\mathbf{b}_p + (\tilde{\mathbf{E}}\mathbf{b}_p)^T\mathbf{C}^{-1}\mathbf{b}_p\right] = \bar{d}_M^2\det(\mathbf{I} - \nu\mathbf{E})^2 - \mathbf{b}_p^T\mathbf{C}^{-1}\mathbf{b}_p$$

Solving this polynomial typically yields up to four real roots, representing different local stationary points (minima and maxima) of the optimization landscape.

The root ν yielding the minimum objective function is selected. Once ν is established, the final B-plane position is deterministically known, and the optimal initial costates required to propagate the fully nonlinear dynamics are recovered analytically [16]:

$$\lambda_{r0} = \mathbf{D}^{-1}\left[(\mathbf{I} - 2\nu\mathbf{D}\mathbf{B}^{-1}\mathbf{R}_{2D}^T\mathbf{C}^{-1}\mathbf{R}_{2D})^{-1}(-2\nu\mathbf{D}\mathbf{B}^{-1}\mathbf{R}_{2D}^T\mathbf{C}^{-1}\mathbf{R}_{2D}\mathbf{r}_s + \mathbf{r}_p) - \mathbf{r}_p\right]$$

$$\lambda_{v0} = -\mathbf{\Phi}_{44}^{-1}\mathbf{\Phi}_{43}\lambda_{r0}$$

2. Miss Distance Constraint:

When a mission operator dictates physical clearance rather than a statistical threshold, the EOCP must be reformulated to enforce a precise physical MD, d, at the Time of Closest Approach. The target constraint alters the boundary formulation to exactly target the Euclidean norm of the final relative position mapped to the B-plane:

$$||\mathbf{R}_{b,2D}(\mathbf{r}_f - \mathbf{r}_s)|| = \bar{d}$$

Following the exact same linearization protocol involving the STM and the submatrices $\mathbf{B}$ and $\mathbf{D}$, the transversality condition updates dynamically based on the Euclidean gradient instead of the Mahalanobis one. This updates the coupled constraint system to:

$$\begin{cases} \mathbf{R}_{b,2D}(\mathbf{r}_f - \mathbf{r}_p) = 2\nu\mathbf{R}_{b,2D}\mathbf{D}\mathbf{B}^{-1}\mathbf{R}_{b,2D}^T\mathbf{R}_{b,2D}(\mathbf{r}_f - \mathbf{r}_s) \\ ||\mathbf{R}_{b,2D}(\mathbf{r}_f - \mathbf{r}_s)|| = \bar{d} \end{cases}$$

By redefining the characteristic mapping matrix[16] without the covariance inverse as:

$$\mathbf{E} = 2\mathbf{R}_{b,2D}\mathbf{D}\mathbf{B}^{-1}\mathbf{R}_{b,2D}^T$$

The final B-plane position simplifies mathematically to $\mathbf{b}_f = (\mathbf{I}_{2\times 2} - \nu\mathbf{E})^{-1}\mathbf{b}_p$. When substituted into the Euclidean distance constraint, the resulting characteristic equation remains a fourth-order polynomial in ν [16]:

$$\nu^2(\tilde{\mathbf{E}}\mathbf{b}_p)^T(\tilde{\mathbf{E}}\mathbf{b}_p) - \nu\mathbf{b}_p^T\tilde{\mathbf{E}}\mathbf{b}_p - \nu(\tilde{\mathbf{E}}\mathbf{b}_p)^T\mathbf{b}_p = \det(\mathbf{I} - \nu\mathbf{E})^2\bar{d}^2 - \mathbf{b}_p^T\mathbf{b}_p$$

After isolating the optimal root for ν, the initial boundary value costates λ_{r0} and λ_{v0} are retrieved directly through algebraic manipulation [16]:

$$\lambda_{r0} = \mathbf{D}^{-1}\left[(\mathbf{I} - 2\nu\mathbf{D}\mathbf{B}^{-1}\mathbf{R}_{2D}^T\mathbf{R}_{2D})^{-1}(-2\nu\mathbf{D}\mathbf{B}^{-1}\mathbf{R}_{2D}^T\mathbf{R}_{2D}\mathbf{r}_s + \mathbf{r}_p) - \mathbf{r}_p\right]$$

$$\lambda_{v0} = -\mathbf{\Phi}_{44}^{-1}\mathbf{\Phi}_{43}\lambda_{r0}$$

These derived initial costates subsequently permit the forward numerical integration of the exact nonlinear Euler-Lagrange equations, effectively bypassing the need for computationally sensitive boundary value solvers while still satisfying the mission's prescribed MD.

3.3.2. Fuel Optimal Control Design

To strictly minimize the utilized propellant mass under realistic hardware constraints, the optimization must be formulated as a Fuel-Optimal Control Problem (FOCP), which incorporates strict bounds on the thrust magnitude and inherently yields a discontinuous, bang-bang control profile. The standard mathematical approach to the FOCP aims to minimize the overall control effort, defined by the performance index:

$$J = \lambda_{v,th} \int_{t_0}^{t_f} a_{max} \epsilon \, dt$$

where $\lambda_{v,th}$ is a positive scaling factor introduced to facilitate the formulation.

To handle the strict physical bounds on the throttle factor $\epsilon \in [0, 1]$, the classical variational approach transitions to PMP.

We formulate the system's Hamiltonian by adjoining the Cartesian equations of motion to the cost function integrand using the position and velocity costate vectors (λ_r and λ_v):

$$H = \lambda_{v,th} a_{max} \epsilon + \lambda_r^T \mathbf{v} + \lambda_v^T \left(-\frac{\mu}{r^3} \mathbf{r} + \epsilon a_{max} \alpha \right)$$

According to the PMP, the optimal control sequence must globally minimize this Hamiltonian at every instant in time. First, minimizing with respect to the thrust pointing direction α (where $||\alpha|| = 1$) dictates that the thrust vector must strictly oppose the velocity costate vector:

$$\alpha^* = -\frac{\lambda_v}{||\lambda_v||}$$

Substituting this optimal direction back into the Hamiltonian allows us to isolate the terms dependent on the throttle fraction ϵ:

$$H = a_{max} \epsilon \big(\lambda_{v,th} - ||\lambda_v|| \big) + \lambda_r^T \mathbf{v} - \lambda_v^T \frac{\mu}{r^3} \mathbf{r}$$

This linear dependency on the control magnitude introduces the fundamental *switching function*, $S(t)$, formally defined as the term multiplying the control variable:

$$S(t) = \lambda_{v,th} - ||\lambda_v(t)||$$

Because the Hamiltonian must be minimized and ϵ is strictly bounded, the optimal throttle factor ϵ^* is entirely dictated by the sign of the switching function. This mathematically

proves the necessity of the fundamental bang-bang optimal control law:

$$\epsilon^* = \begin{cases} 1 & \text{if } S(t) < 0 \implies ||\lambda_v|| > \lambda_{v,th} \quad \text{(Maximum Thrust)} \\ 0 & \text{if } S(t) > 0 \implies ||\lambda_v|| < \lambda_{v,th} \quad \text{(Coast)} \end{cases}$$

The physical transitions between active thrusting and coasting phases occur precisely at the singular points where the switching function evaluates to zero ($S(t) = 0$). Consequently, the engine ignition time (t_{on}) and cutoff time (t_{off}) are mathematically defined as the temporal roots of this switching function. Identifying these exact switching bounds is therefore equivalent to finding the precise epochs where the magnitude of the velocity costate intersects the scalar threshold $\lambda_{v,th}$.

Because discontinuous pulse functions are notoriously difficult to resolve using standard numerical TPBVP solvers (often requiring imprecise hyperbolic tangent smoothing functions), a novel **Differential Correction (DC) algorithm** algorithm is proposed in Figure 3.2.

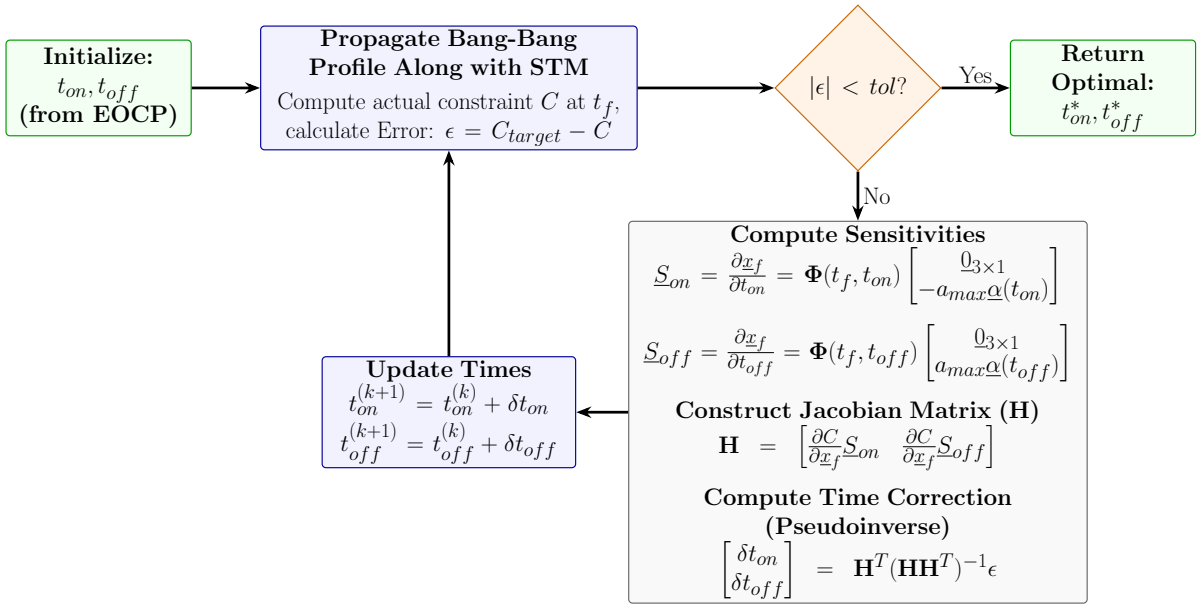

Figure 3.2: Simplified block diagram of the Differential Correction method.

This method abandons the integration of costate differential equations during the correction phase entirely. Instead, it relies on a fundamental numerical observation derived from the optimal control analysis made by De Vittori et. al. [16]: the time evolution of the velocity costate vector, $\lambda_v(t)$, remains remarkably similar between the energy-optimal and the true fuel-optimal solutions as displayed in Figure 3.3.

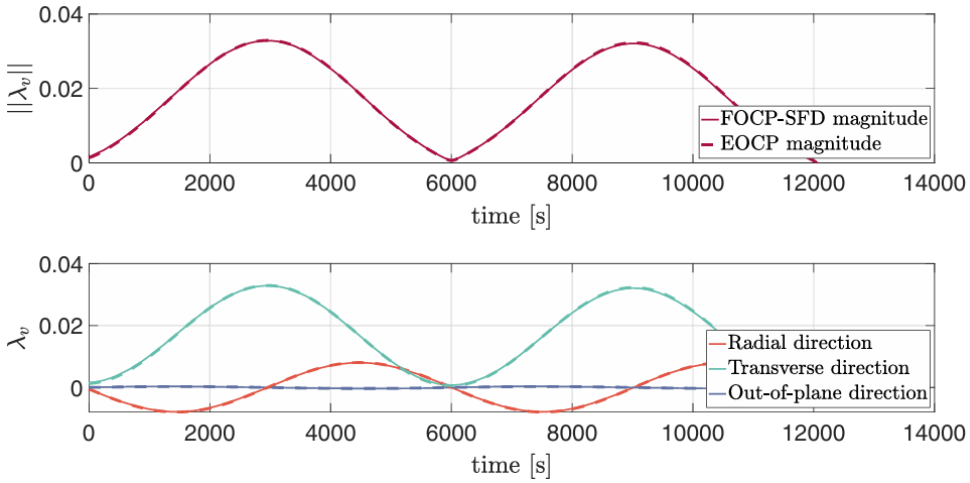

Figure 3.3: Velocity costate profile comparison. The energy-optimal CAM (dashed lines) closely mirrors the true fuel-optimal control (solid lines) initiated two orbits before TCA [16]

By freezing the thrust direction to the analytical energy-optimal vector, the optimization problem simplifies strictly to identifying the correct temporal switching structure—namely, the precise engine ignition time (t_{on}) and the engine cutoff time (t_{off}) of a rigid bang-bang profile.

To iteratively satisfy the terminal safety requirements via differential correction, it is necessary to determine exactly how minute adjustments to the engine switching times (δt_{on}, δt_{off}) influence the scalar constraint value C evaluated at TCA. Because the constraint is evaluated exclusively based on the spatial proximity at the conjunction epoch, the total constraint variation δC must first be expressed in terms of the final state variation, $\delta \mathbf{x}_f$:

$$\delta C = \frac{\partial C}{\partial \mathbf{x}_f} \delta \mathbf{x}_f$$

Because the thrust pointing direction is frozen, this variation of the final state is driven entirely by perturbations in the temporal switching bounds. Applying the chain rule, and explicitly defining the kinematic state sensitivities as $\mathbf{S}_{on} = \frac{\partial \mathbf{x}_f}{\partial t_{on}}$ and $\mathbf{S}_{off} = \frac{\partial \mathbf{x}_f}{\partial t_{off}}$, we expand $\delta \mathbf{x}_f$:

$$\delta \mathbf{x}_f = \frac{\partial \mathbf{x}_f}{\partial t_{on}} \delta t_{on} + \frac{\partial \mathbf{x}_f}{\partial t_{off}} \delta t_{off} = \mathbf{S}_{on} \delta t_{on} + \mathbf{S}_{off} \delta t_{off}$$

Substituting this expanded state variation back into the initial constraint differential yields:

$$\begin{aligned}\delta C &= \frac{\partial C}{\partial \mathbf{x}_f}\left(\mathbf{S}_{on}\delta t_{on} + \mathbf{S}_{off}\delta t_{off}\right) \\ &= \left(\frac{\partial C}{\partial \mathbf{x}_f}\mathbf{S}_{on}\right)\delta t_{on} + \left(\frac{\partial C}{\partial \mathbf{x}_f}\mathbf{S}_{off}\right)\delta t_{off}\end{aligned} \tag{3.1}$$

By extracting the coupled sensitivity terms, we define the central 1×2 Jacobian matrix, $\mathbf{H} = \begin{bmatrix}\frac{\partial C}{\partial \mathbf{x}_f}\mathbf{S}_{on} & \frac{\partial C}{\partial \mathbf{x}_f}\mathbf{S}_{off}\end{bmatrix}$. This allows the entire differential relationship to be elegantly expressed as:

$$\delta C = \mathbf{H}\begin{bmatrix}\delta t_{on} \\ \delta t_{off}\end{bmatrix}$$

To systematically compute these sensitivities ($\mathbf{S}_{on}$ and $\mathbf{S}_{off}$), the STM, $\mathbf{\Phi}(t, t_0)$, is utilized. The STM is a fundamental mathematical tool that maps an instantaneous state perturbation at an initial epoch forward to a final epoch, visualized in Figure 3.4 through the linearized flow of the system.

$$\mathbf{\Phi}(t, t_0) = \frac{\partial \mathbf{x}}{\partial \mathbf{x}_0} = \begin{bmatrix}\frac{\partial \mathbf{r}}{\partial \mathbf{r}_0} & \frac{\partial \mathbf{r}}{\partial \mathbf{v}_0} \\ \frac{\partial \mathbf{v}}{\partial \mathbf{r}_0} & \frac{\partial \mathbf{v}}{\partial \mathbf{v}_0}\end{bmatrix}$$

Its dynamics are governed by the linear matrix differential equation:

$$\dot{\mathbf{\Phi}}(t, t_0) = \mathbf{A}(t)\mathbf{\Phi}(t, t_0), \quad \mathbf{\Phi}(t_0, t_0) = \mathbf{I}_{6\times 6}$$

where $\mathbf{A}(t) = \frac{\partial \mathbf{f}}{\partial \mathbf{x}}$ is the Jacobian matrix of the two-body dynamics evaluated along the reference trajectory, as computed in Eq. (3.2). By integrating this matrix differential equation, the STM encapsulates how variations in position and velocity naturally evolve over time under the influence of the Earth's gravity gradient.

Using the standard Two-Body Problem (2BP) STM during active thrusting is a valid approximation. Because low-thrust propulsion only slightly alters the nominal trajectory during the short maneuver window, the system Jacobian can rely strictly on conservative gravitational terms.
Applying this to the classical 2BP, the state vector derivative is defined by:

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}) = \left[v_x,\ v_y,\ v_z,\ -\mu x r^{-3},\ -\mu y r^{-3},\ -\mu z r^{-3}\right]^T = \left[\mathbf{v},\ \mathbf{a}_{grav}\right]^T, \quad r = \sqrt{x^2 + y^2 + z^2}$$

The system matrix $\mathbf{A}(t)$ is the Jacobian of this function:

$$\mathbf{A}(t) = \frac{\partial \mathbf{f}}{\partial \mathbf{x}} = \begin{bmatrix} \mathbf{0}_{3\times3} & \mathbf{I}_{3\times3} \\ \frac{3\mu}{r^5}\begin{bmatrix} x^2 & xy & xz \\ xy & y^2 & yz \\ xz & yz & z^2 \end{bmatrix} - \frac{\mu}{r^3}\mathbf{I}_{3\times3} & \mathbf{0}_{3\times3} \end{bmatrix} = \begin{bmatrix} \frac{\partial \mathbf{v}}{\partial \mathbf{r}} & \frac{\partial \mathbf{v}}{\partial \mathbf{v}} \\ \frac{\partial \mathbf{a}_{grav}}{\partial \mathbf{r}} & \frac{\partial \mathbf{a}_{grav}}{\partial \mathbf{v}} \end{bmatrix} \tag{3.2}$$

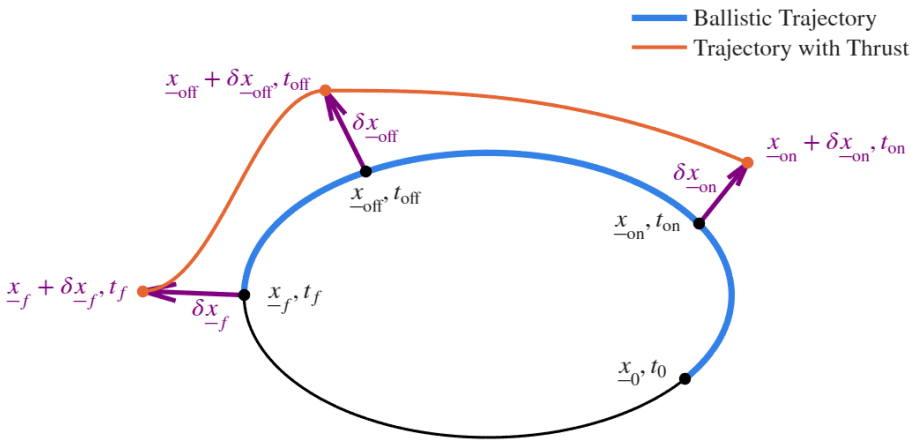

Figure 3.4: Qualitative visualization of the STM propagating an instantaneous perturbation induced by the propulsion unit forward to the terminal state.

With the system Jacobian established and the STM dynamics fully defined, the process transitions to evaluating the physical effects of modifying the maneuver's temporal bounds. This evaluation maps perturbations in the switching limits directly to the terminal state, a procedure analyzed separately for the ignition and cutoff times in the following two parts.

1. Sensitivity to Ignition Time (t_{on}):

Delaying the engine ignition by a small positive temporal increment δt effectively removes a slice of thrust that would have otherwise occurred at the beginning of the burn arc. As visualized in Figure 3.5, this creates a deficit in the total integrated acceleration.

$$\delta \mathbf{v}_{on} = \int_{t_{on}}^{t_{on}+\delta t} -\mathbf{a}_{burn}\, dt = \int_{t_{on}}^{t_{on}+\delta t} -a_{max}\alpha\, dt \approx -a_{max}\alpha(t_{on})\delta t$$

Because spatial position cannot change instantaneously, the instantaneous 6-dimensional state variation at ignition is purely kinematic.
Mapping this initial perturbation forward through the orbital dynamics to the final conjunction epoch t_f explicitly requires the STM evaluated from t_{on} to t_f, denoted as

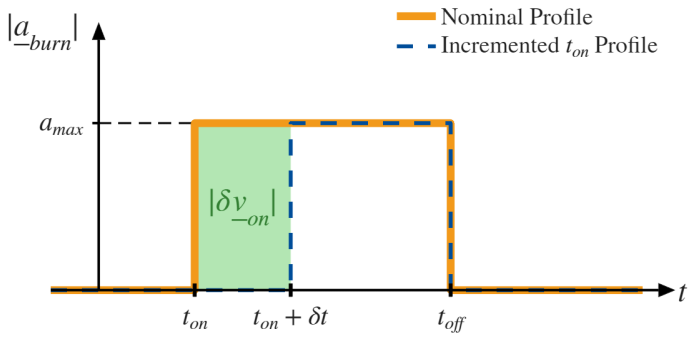

Figure 3.5: Effect of incrementing the ignition time t_{on}, visually demonstrating the removal of a slice of the active thrust profile.

$\mathbf{\Phi}(t_f, t_{on})$. This yields the 6×1 state sensitivity to the ignition time:

$$\mathbf{S}_{on} = \frac{\partial \mathbf{x}_f}{\partial t_{on}} = \mathbf{\Phi}(t_f, t_{on}) \begin{bmatrix} \mathbf{0}_{3\times 1} \\ -a_{max}\alpha(t_{on}) \end{bmatrix}$$

2. Sensitivity to Cutoff Time (t_{off}):

Conversely, extending the engine cutoff time by a small positive temporal increment δt adds an additional slice of thrust at the end of the burn arc, as shown in Figure 3.6.

$$\delta \mathbf{v}_{off} = \int_{t_{off}}^{t_{off}+\delta t} \mathbf{a}_{burn}\, dt = \int_{t_{off}}^{t_{off}+\delta t} a_{max}\alpha\, dt \approx a_{max}\alpha(t_{off})\delta t$$

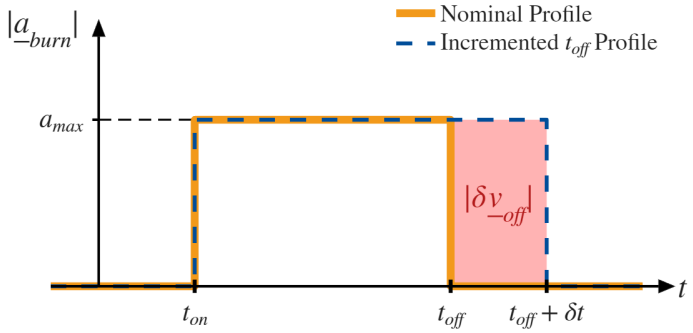

Figure 3.6: Effect of incrementing the cutoff time t_{off}, visually demonstrating the addition of a slice to the active thrust profile.

Mapping this positive velocity perturbation forward to the conjunction epoch via the appropriate STM yields the 6×1 state sensitivity to the cutoff time:

$$\mathbf{S}_{off} = \frac{\partial \mathbf{x}_f}{\partial t_{off}} = \mathbf{\Phi}(t_f, t_{off}) \begin{bmatrix} \mathbf{0}_{3\times 1} \\ a_{max}\alpha(t_{off}) \end{bmatrix}$$

These sensitivity vectors $\mathbf{S}_{on}$ and $\mathbf{S}_{off}$ form the core mapping mechanics. They are projected onto the specific gradient ($\frac{\partial C}{\partial \mathbf{x}_f}$) of the required constraint to complete the formulation of Eq. (3.1). Depending on the operational requirements, the target constraint evaluates either a deterministic MD or a statistical boundary, as detailed in the subsequent two formulations.

1. Miss Distance Constraint:

When mission operations dictate a strict geometric clearance, the target constraint C evaluates the actual MD at TCA. Projected onto the B-plane, the actual MD d achieved by the current switching times is the Euclidean norm of the relative position vector $\mathbf{b}_f = \mathbf{R}_{b,2D}(\mathbf{r}_f - \mathbf{r}_s)$:

$$C = d = ||\mathbf{b}_f|| = \sqrt{\mathbf{b}_f^T \mathbf{b}_f}$$

Applying the chain rule, the sensitivity with respect to the final spatial position is evaluated as:

$$\frac{\partial d}{\partial \mathbf{r}_f} = \frac{\partial d}{\partial \mathbf{b}_f}\frac{\partial \mathbf{b}_f}{\partial \mathbf{r}_f} = \frac{\partial \sqrt{\mathbf{b}_f^T \mathbf{b}_f}}{\partial \mathbf{b}_f}\frac{\partial \mathbf{b}_f}{\partial \mathbf{r}_f} = \frac{\mathbf{b}_f^T}{||\mathbf{b}_f||}\mathbf{R}_{b,2D} = \frac{1}{d}\mathbf{b}_f^T \mathbf{R}_{b,2D}$$

As this geometric metric has no direct velocity dependence, $\frac{\partial d}{\partial \mathbf{v}_f} = \mathbf{0}_{1\times 3}$. This formulation yields the 1×6 terminal state sensitivity gradient:

$$\frac{\partial d}{\partial \mathbf{x}_f} = \begin{bmatrix} \frac{1}{d}\mathbf{b}_f^T \mathbf{R}_{b,2D} & \mathbf{0}_{1\times 3} \end{bmatrix}$$

By defining the 1×2 Jacobian matrix $\mathbf{H} = \begin{bmatrix} \frac{\partial d}{\partial \mathbf{x}_f}\mathbf{S}_{on} & \frac{\partial d}{\partial \mathbf{x}_f}\mathbf{S}_{off} \end{bmatrix}$, the required temporal corrections to target a specific $\bar{d}$ are evaluated utilizing the Moore-Penrose pseudoinverse.

2. Squared Mahalanobis Distance Constraint:

If the operational safety requirement is framed probabilistically, the CAM design algorithm must target the SMD (d_M^2) rather than the deterministic physical MD. The core architecture of the differential correction algorithm remains identical; only the terminal sensitivity gradient is adjusted. The parameter d_M^2 is defined as a statistical quadratic form involving the combined B-plane covariance matrix $\mathbf{C}$:

$$C = d_M^2 = \mathbf{b}_f^T \mathbf{C}^{-1} \mathbf{b}_f \quad \text{where} \quad \mathbf{b}_f = \mathbf{R}_{b,2D}(\mathbf{r}_f - \mathbf{r}_s)$$

Applying the chain rule, the sensitivity of d_M^2 with respect to the final spatial position in the ECI frame is:

$$\frac{\partial d_M^2}{\partial \mathbf{r}_f} = \frac{\partial d_M^2}{\partial \mathbf{b}_f}\frac{\partial \mathbf{b}_f}{\partial \mathbf{r}_f} = 2\mathbf{b}_f^T\mathbf{C}^{-1}\mathbf{R}_{b,2D}$$

As before, this statistical metric has no direct velocity dependence, so $\frac{\partial d_M^2}{\partial \mathbf{v}_f} = \mathbf{0}_{1\times 3}$. This formulation yields the updated 1×6 terminal state sensitivity gradient for the Jacobian matrix $\mathbf{H}$:

$$\frac{\partial d_M^2}{\partial \mathbf{x}_f} = \begin{bmatrix} 2\mathbf{b}_f^T\mathbf{C}^{-1}\mathbf{R}_{b,2D} & \mathbf{0}_{1\times 3} \end{bmatrix}$$

By defining the 1×2 Jacobian matrix $\mathbf{H} = \begin{bmatrix} \frac{\partial d_M^2}{\partial \mathbf{x}_f}\mathbf{S}_{on} & \frac{\partial d_M^2}{\partial \mathbf{x}_f}\mathbf{S}_{off} \end{bmatrix}$, the identical pseudoinverse logic is deployed to iteratively correct the switching times. This elegant substitution ensures that the Differential Correction architecture seamlessly transitions between targeting simple circular exclusion zones and complex, covariance-scaled elliptical boundaries.

3.4. Algorithm Workflow: Energy-Optimal Initialization and Differential Correction Loop

To overcome the severe numerical sensitivities associated with solving the discontinuous Fuel-Optimal Control Problem (FOCP) via traditional Two-Point Boundary Value Problem (TPBVP) solvers, this chapter proposes a highly efficient two-phase algorithm. This architecture seamlessly bridges the computational speed of analytical energy-optimal formulations with the strict mass-efficiency of a numerical differential correction loop. The complete methodological architecture is visually summarized in the block diagram in Figure 3.7 and formally detailed in Algorithm 3.2.

3.4.1. Phase 1: Energy-Optimal Initialization

The first phase is designed to circumvent the need for arbitrary initial guesses by providing an analytically backed, highly accurate estimation of the optimal thrust pointing direction and the initial temporal switching bounds. This initialization phase executes the following sequential steps to extract the operational data and generate the initial conditions.

1. Data Extraction from Samples: The procedure begins by extracting the necessary Conjunction Data Message (CDM) parameters evaluated at the Time of Closest Approach (TCA), denoted as t_f. The extracted data includes the primary and secondary spacecraft states ($\mathbf{x}_p(t_f)$, $\mathbf{x}_s(t_f)$), their respective covariance matrices ($\mathbf{C}_p$, $\mathbf{C}_s$), and the combined hardbody radius (s_A).

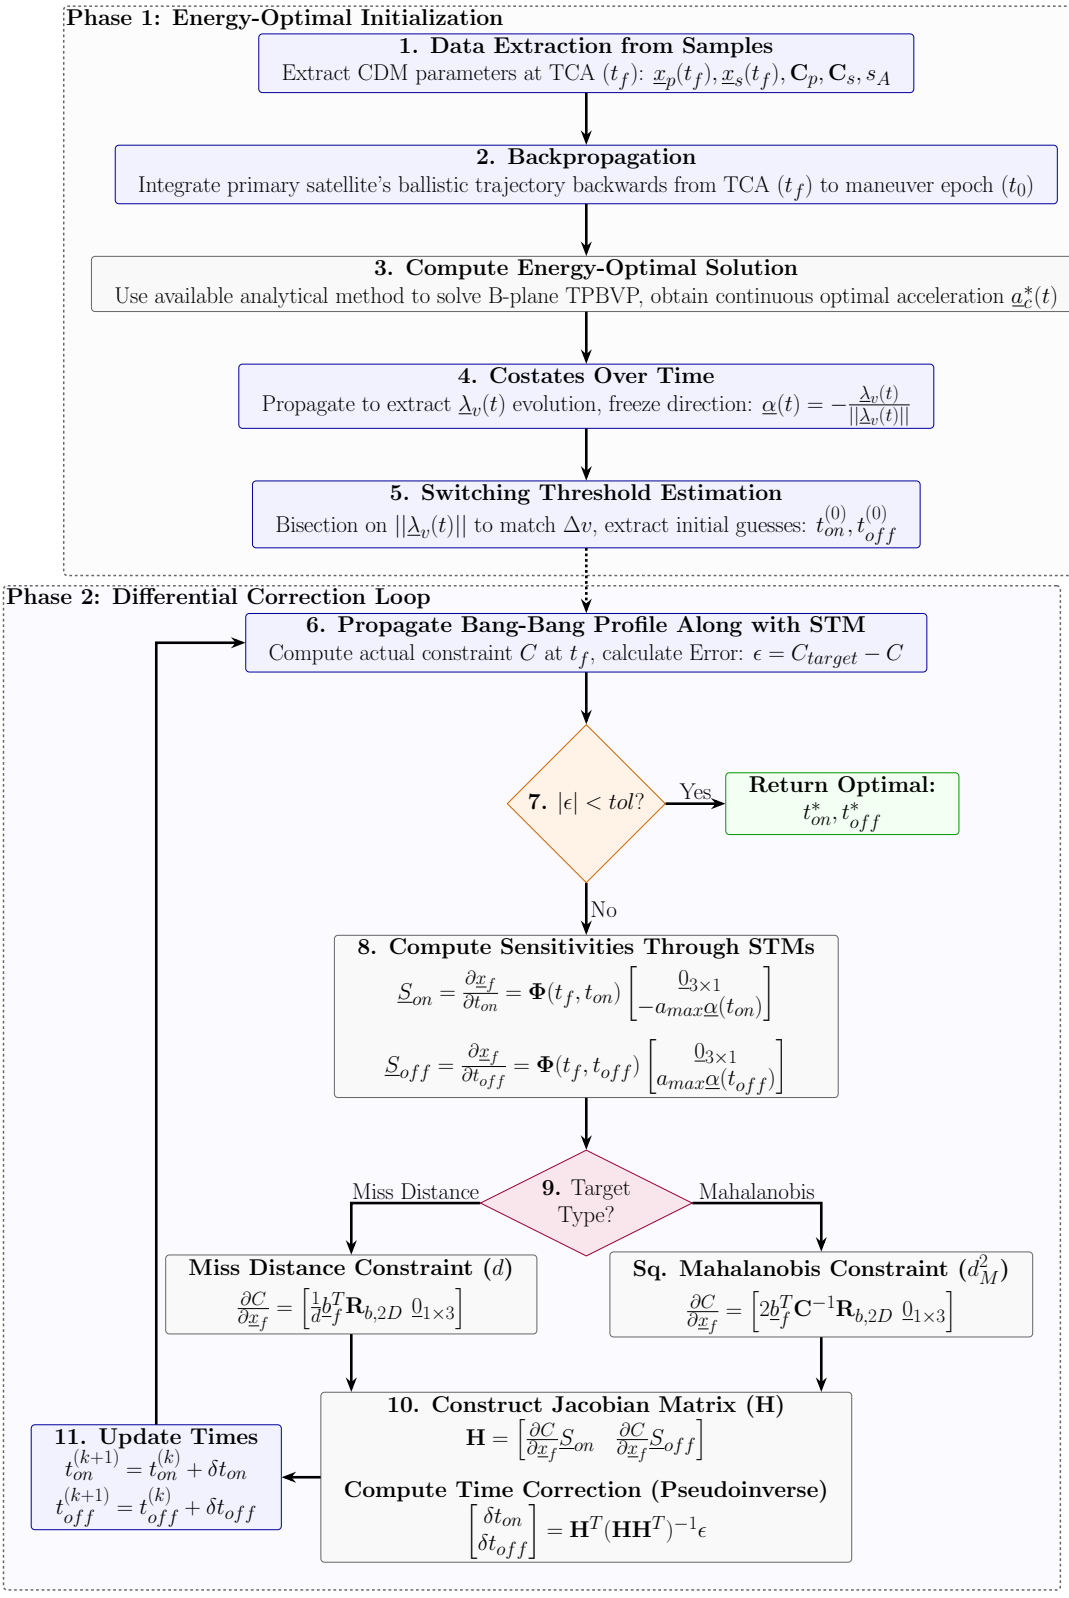

Figure 3.7: Complete workflow of the proposed methodology, highlighting the transition from the analytical energy-optimal initialization (Phase 1) to the numerical differential correction loop (Phase 2).

2. Backpropagation: The nominal trajectories of the primary satellite is integrated backwards following ballistic 2BP dynamics in time from the collision epoch (t_f) to the designated maneuver initiation epoch (t_0).

3. Compute Energy-Optimal Solution: An analytical formulation is utilized to solve the unconstrained Energy-Optimal Control Problem (EOCP) projected on the B-plane. This yields a continuous, analytically defined optimal state-costate solution.

4. Costates Over Time: By propagating the system, the evolution of the velocity costate vector $\lambda_v(t)$ is extracted. Leveraging the geometric duality between energy- and fuel-optimal solutions, the near-optimal thrust pointing direction $\alpha(t)$ is formally frozen for the remainder of the optimization process:

$$\alpha(t) = -\frac{\lambda_v(t)}{||\lambda_v(t)||}$$

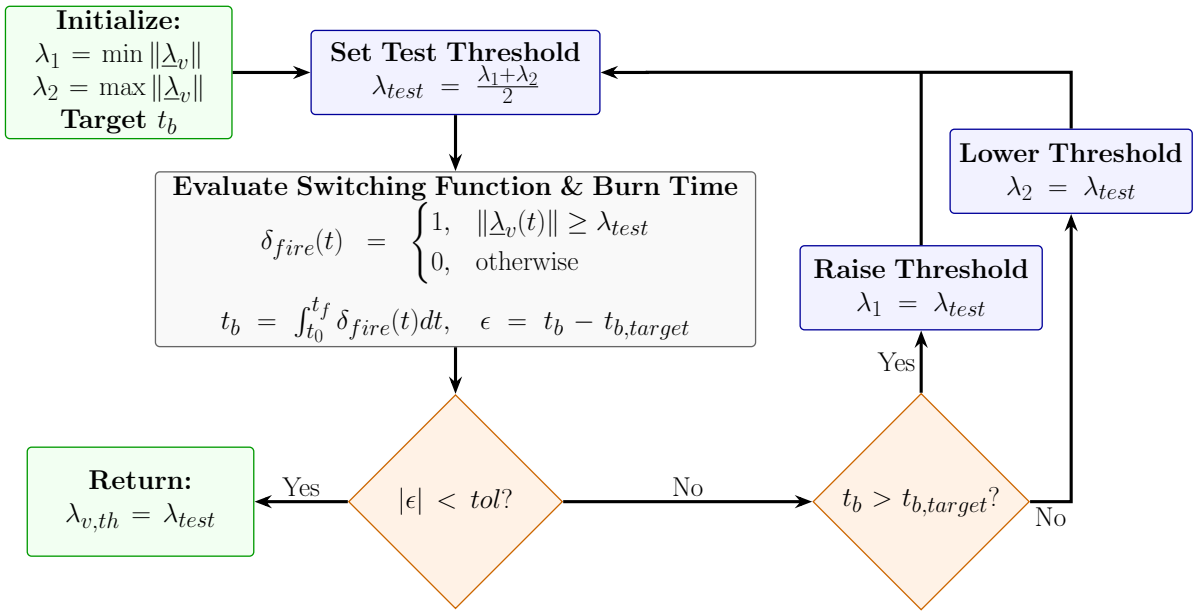

Figure 3.8: Block diagram of the Bisection Method for Velocity Costate Threshold

5. Switching Threshold Estimation: Because the engine must operate on a strict bang-bang principle, a scalar switching threshold is required. The rapid bisection method detailed in Algorithm 3.1 and Figure 3.8 is applied to the velocity costate magnitude, $||\lambda_v(t)||$, to find the exact threshold that matches the analytically available Δv. Because the control acceleration is constrained to a constant maximum amplitude (a_{max}), as mathematically defined starting from Eq. (2.1), the equivalent continuous burning time (t_b)

can be computed directly:

$$t_b = \frac{\Delta v}{a_{max}} = \int_{t_0}^{t_f} ||\lambda_v(t)||\, dt$$

The temporal roots of this intersection provide the precise initial guesses for the engine ignition and cutoff times:

$$||\lambda_v(t)|| - \lambda_{th} = 0 \implies \text{Initial Guesses: } \{t_{on}^{(0)}, t_{off}^{(0)}\}$$

Algorithm 3.1 Bisection Method for Velocity Costate Threshold

Input: Time array t, energy-optimal velocity costate history $\lambda_v(t)$, target equivalent burning time t_b, convergence tolerance tol, maximum iterations N_{max}
Output: Velocity costate threshold $\lambda_{v,th}$
$\lambda_1 \leftarrow \min(||\lambda_v(t)||)$
$\lambda_2 \leftarrow \max(||\lambda_v(t)||)$
$best_\lambda \leftarrow \frac{\lambda_1+\lambda_2}{2}$
$min_diff \leftarrow \infty$
for $k = 1, 2, \ldots, N_{max}$ **do**
  $\lambda_{test} \leftarrow \frac{\lambda_1+\lambda_2}{2}$
  $\delta_{fire}(t) \leftarrow \begin{cases} 1, & \text{if } ||\lambda_v(t)|| \geq \lambda_{test} \\ 0, & \text{otherwise} \end{cases}$
  $t_{burn} \leftarrow \int_{t_0}^{t_f} \delta_{fire}(t)\, dt$
  $\varepsilon \leftarrow t_{burn} - t_b$
  if $|\varepsilon| < min_diff$ **then**
    $min_diff \leftarrow |\varepsilon|$
    $best_\lambda \leftarrow \lambda_{test}$
  end if
  if $|\varepsilon| < tol$ **then**
    return λ_{test} *(Convergence condition met)*
  end if
  if $t_{burn} > t_b$ **then**
    $\lambda_1 \leftarrow \lambda_{test}$ *(Firing too long → raise threshold)*
  else
    $\lambda_2 \leftarrow \lambda_{test}$ *(Firing too short → lower threshold)*
  end if
end for
return $best_\lambda$ *(Fallback if max iterations reached)*

3.4.2. Phase 2: Differential Correction Loop

With the thrust direction $\alpha(t)$ frozen and the initial switching bounds $\{t_{on}^{(k)}, t_{off}^{(k)}\}$ established, the second phase executes a Differential Correction (DC) loop to iteratively perturb these burning times until the terminal safety constraints are precisely satisfied. This correction loop iterates through the following procedural steps to refine the switching bounds.

6. Propagate Bang-Bang State: The complete coupled system—comprising the spacecraft state vector $\mathbf{x} = [\mathbf{r}^T, \mathbf{v}^T]^T$ and the 6×6 STM $\mathbf{\Phi}(t, t_0)$—is simultaneously forward-propagated from t_0 to t_f. The whole 42-dimensional differential dynamics system is governed by:

$$\frac{d}{dt}\begin{bmatrix} \mathbf{r}(t) \\ \mathbf{v}(t) \\ \mathbf{\Phi}(t,t_0) \end{bmatrix} = \begin{bmatrix} \mathbf{v}(t) \\ -\frac{\mu}{r^3(t)}\mathbf{r}(t) + \epsilon(t) a_{max}\alpha(t) \\ \mathbf{A}(t)\mathbf{\Phi}(t,t_0) \end{bmatrix}$$

where $\epsilon(t) = 1$ for $t \in [t_{on}^{(k)}, t_{off}^{(k)}]$ and 0 otherwise. The system matrix $\mathbf{A}(t)$ evaluates the Jacobian of the dynamics along the trajectory:

$$\mathbf{A}(t) = \begin{bmatrix} \mathbf{0}_{3\times3} & \mathbf{I}_{3\times3} \\ \frac{\mu}{r^3}\left(\frac{3\mathbf{r}\mathbf{r}^T}{r^2} - \mathbf{I}_{3\times3}\right) & \mathbf{0}_{3\times3} \end{bmatrix}$$

At t_f, the constraint function C is evaluated, and the residual error is computed as:

$$\varepsilon = C_{target} - C$$

7. Decision Check: The absolute error $|\epsilon|$ is evaluated against a defined tolerance (tol). If $|\epsilon| < tol$, the loop terminates successfully, returning the optimal times t_{on}^* and t_{off}^*. Otherwise, the algorithm proceeds to compute corrections.

8. Compute Sensitivities Through STMs: The variations of the terminal state $\mathbf{x}_f$ with respect to the temporal switching limits are computed utilizing the STM.

A delay in ignition translates to less thrust, hence the negative sign, while a delay in cutoff translates to more thrust:

$$\mathbf{S}_{on} = \frac{\partial \mathbf{x}_f}{\partial t_{on}} = \mathbf{\Phi}(t_f, t_{on}) \begin{bmatrix} \mathbf{0}_{3\times1} \\ -a_{max}\alpha(t_{on}) \end{bmatrix} \tag{3.3}$$

$$\mathbf{S}_{off} = \frac{\partial \mathbf{x}_f}{\partial t_{off}} = \mathbf{\Phi}(t_f, t_{off}) \begin{bmatrix} \mathbf{0}_{3\times1} \\ a_{max}\alpha(t_{off}) \end{bmatrix} \tag{3.4}$$

9. Target Type Decision & Gradient Calculation: The analytical gradient of the terminal constraint function, denoted as $\frac{\partial C}{\partial \mathbf{x}_f}$, forms a critical 1×6 vector that bridges the fully coupled 3D inertial dynamics to the localized 2D safety metrics on the B-plane. Because the safety criteria—whether geometric or statistical—are evaluated exclusively based on spatial proximity at the Time of Closest Approach (TCA), there is no direct dependency on the terminal velocity vector. Consequently, the velocity components of the gradient are strictly padded with zeros ($\mathbf{0}_{1\times3}$). The position components are dynamically determined based on the specific operational safety criteria chosen by the mission designer:

- **Miss Distance Constraint (d):** When targeting a strict physical separation distance, the gradient evaluates the Euclidean norm sensitivity:

$$\frac{\partial C}{\partial \mathbf{x}_f} = \begin{bmatrix} \frac{1}{d}\mathbf{b}_f^T\mathbf{R}_{b,2D} & \mathbf{0}_{1\times3} \end{bmatrix}$$

- **Squared Mahalanobis Constraint (d_M^2):** When the mission dictates a probabilistic risk threshold, the gradient evaluates the sensitivity of the covariance-scaled exclusion ellipse:

$$\frac{\partial C}{\partial \mathbf{x}_f} = \begin{bmatrix} 2\mathbf{b}_f^T\mathbf{C}^{-1}\mathbf{R}_{b,2D} & \mathbf{0}_{1\times3} \end{bmatrix}$$

10. Construct Jacobian: With both the terminal constraint gradient and the STM-propagated temporal sensitivities ($\mathbf{S}_{on}$ and $\mathbf{S}_{off}$) established, they are multiplied via the chain rule to formulate the 1×2 Jacobian matrix $\mathbf{H}$. This matrix serves as the fundamental engine of the differential correction loop. It mathematically collapses the 6-dimensional state space and linearly connects minute temporal perturbations in the switching bounds directly to the resulting scalar error on the B-plane.

$$\mathbf{H} = \begin{bmatrix} \frac{\partial C}{\partial \mathbf{x}_f}\mathbf{S}_{on} & \frac{\partial C}{\partial \mathbf{x}_f}\mathbf{S}_{off} \end{bmatrix}$$

11. Update Times: Because the algorithm adjusts two independent control variables (t_{on} and t_{off}) to satisfy a single scalar constraint (C), the system is mathematically underdetermined. To iteratively drive the current constraint value $C(\mathbf{x}_f)$ toward the desired target C_{target}, a first-order Taylor expansion is applied:

$$C_{target} \approx C(\mathbf{x}_f) + \delta C$$

By rearranging this expansion, it becomes clear that the required differential constraint

correction δC is exactly equivalent to the evaluated residual error ϵ:

$$\delta C = C_{target} - C(\mathbf{x}_f) = \epsilon$$

Algorithm 3.2 Differential Correction for Fuel-Optimal Switching Times

Initial instructions: Energy-optimal thrust direction $\alpha(t)$, max acceleration a_{max}, target constraint C_{target} ($\bar{d}$ or $\bar{d}_M^2$), tolerance tol, max iterations N_{max}
PHASE 1: Initialize via Bisection Method
$\lambda_{v,th} \leftarrow$ Run Algorithm 1 (Bisection Method)
Extract initial t_{on} and t_{off} such that $||\lambda_v(t)|| = \lambda_{v,th}$
PHASE 2: Differential Correction Loop
for $k = 1, 2, \ldots, N_{max}$ **do**
  Propagate trajectory from t_0 to t_f with $\mathbf{a}_{burn} = a_{max}\alpha$ and simultaneously integrate STM $\mathbf{\Phi}(t, t_0)$
  Compute actual constraint value C (d or d_M^2) at t_f, and error $\varepsilon \leftarrow C_{target} - C$
  if $|\varepsilon| < tol$ **then**
    return t_{on}, t_{off} *(Convergence met)*
  end if
  Extract $\mathbf{\Phi}(t_f, t_{on})$ and $\mathbf{\Phi}(t_f, t_{off})$ from the integrated matrix history
  $\mathbf{S}_{on} \leftarrow \mathbf{\Phi}(t_f, t_{on})\begin{bmatrix}\mathbf{0}_{1\times3}, & -a_{max}\alpha^T(t_{on})\end{bmatrix}^T$
  $\mathbf{S}_{off} \leftarrow \mathbf{\Phi}(t_f, t_{off})\begin{bmatrix}\mathbf{0}_{1\times3}, & a_{max}\alpha^T(t_{off})\end{bmatrix}^T$
  if Target is Miss Distance **then**
    $\nabla C \leftarrow \begin{bmatrix}\frac{1}{d}\mathbf{b}_f^T\mathbf{R}_{b,2D}, & \mathbf{0}_{1\times3}\end{bmatrix}$
  else
    $\nabla C \leftarrow \begin{bmatrix}2\mathbf{b}_f^T\mathbf{C}^{-1}\mathbf{R}_{b,2D}, & \mathbf{0}_{1\times3}\end{bmatrix}$ *(Target is Sq. Mahalanobis)*
  end if
  $\mathbf{H} \leftarrow [\nabla C \cdot \mathbf{S}_{on} \quad \nabla C \cdot \mathbf{S}_{off}]$
  $\begin{bmatrix}\delta t_{on}, & \delta t_{off}\end{bmatrix}^T \leftarrow \mathbf{H}^T(\mathbf{H}\mathbf{H}^T)^{-1}\varepsilon$
  $t_{on} \leftarrow t_{on} + \delta t_{on}$ and $t_{off} \leftarrow t_{off} + \delta t_{off}$
end for
return t_{on}, t_{off} *(Fallback if max iterations reached)*

Substituting this error ϵ back into the linear differential relationship ($\delta C = \mathbf{H}[\delta t_{on}, \delta t_{off}]^T$) explicitly justifies the update law.

To resolve the underdetermined system and ensure the temporal adjustments remain sufficiently small for the linear assumptions of the STM to hold true, a minimum-norm

correction is computed utilizing the Moore-Penrose pseudoinverse:

$$\begin{bmatrix} \delta t_{on} \\ \delta t_{off} \end{bmatrix} = \mathbf{H}^T(\mathbf{H}\mathbf{H}^T)^{-1}\epsilon$$

Once the optimal timing increments δt_{on} and δt_{off} are computed, the current switching times are incrementally updated to form the control sequence for the next iteration:

$$\begin{cases} t_{on}^{(k+1)} = t_{on}^{(k)} + \delta t_{on} \\ t_{off}^{(k+1)} = t_{off}^{(k)} + \delta t_{off} \end{cases}$$

3.5. Summary of This Chapter

This chapter synthesizes the mathematical models and orbital dynamics introduced previously to construct a highly efficient, applied optimal control framework for low-thrust collision avoidance.

1) It begins by detailing the analytical resolution of the Energy-Optimal Control Problem, demonstrating how the complex boundary conditions on the B-plane can be reduced to a polynomial root-finding exercise using STM linearizations.

2) However, recognizing the operational limitations of continuous, unbounded thrust profiles, the formulation transitions to a Fuel-Optimal Control Problem. By evaluating the PMP, it is proven that true fuel-optimality inherently dictates a strict discontinuous bang-bang profile. Because resolving such discontinuous profiles via traditional TPBVP solvers frequently induces severe numerical instability, a novel DC methodology is proposed.

3) The DC approach bypasses costate integration entirely by freezing the thrust pointing direction to the robust energy-optimal solution. Utilizing a rapid bisection method on the velocity costate magnitude, high-fidelity initial estimates for the engine ignition and cutoff times are secured. The algorithm then enters an iterative correction loop, concurrently propagating the spacecraft's trajectory alongside the STM differential equations. This dual propagation allows the algorithm to linearly map variations in the temporal switching bounds directly to the resulting spatial deviations at the conjunction epoch. By strictly defining the terminal gradients for both the MD and the covariance-scaled SMD constraints, the established DC framework is fully equipped to satisfy both deterministic and probabilistic mission safety requirements while explicitly minimizing propellant consumption.

4 | Performance Evaluation of the Low-Thrust Fuel Optimal Differential Correction

4.1. Introduction

This chapter presents the numerical results and the performance evaluation of the proposed Differential Correction (DC) methodology for low-thrust Collision Avoidance Maneuvers (CAMs). To rigorously validate the efficiency, accuracy, and fuel-optimality of the novel approach, it is benchmarked directly against the established state-of-the-art methods recently evaluated in the comprehensive *CAMmary* review [43]. To ensure the reproducibility of the experiments and provide a transparent baseline for the computational time comparisons, all numerical simulations, trajectory propagations, and optimization algorithms presented in this chapter are executed on a single desktop workstation. The hardware and software specifications of the computational environment are detailed below:

- Operating System: Microsoft Windows 11 Pro (64-bit, Version 10.0.26200)
- Processor (CPU): Intel® Core™ i7-10700 @ 2.90 GHz (8 Cores, 16 Logical Processors)
- System Memory (RAM): 16.0 GB
- System Model: LENOVO 90NA001QCD (TianYi510Pro-14IMB)
- Architecture: x64-based PC
- BIOS Version: LENOVO O4HKT30A (2021-03-08)

To rigorously validate the proposed optimal control strategies and the differential correction methodology, a comprehensive set of realistic test cases is utilized. This dataset is sourced from the official Collision Avoidance Challenge organized by the European Space

Agency (ESA) and is publicly available for download[1]. The selected dataset comprises 2,170 individual Conjunction Data Messages (CDMs), each representing a distinct, high-fidelity short-term encounter scenario located within the Low Earth Orbit (LEO) regime. To ensure the optimization algorithms are robustly tested against a diverse and representative spectrum of conjunction geometries and risk profiles, the catalog encompasses encounters with predicted physical miss distances ranging from as close as 8.2 meters up to 2.0 kilometers. Correspondingly, the assessed probabilities of collision within these test cases span a wide operational range, from a baseline safety threshold of 10^{-6} up to a highly critical 0.136.

To ensure a strictly fair and one-to-one comparative analysis, the simulation environment and the hardware constraints are configured to perfectly replicate the conditions established in the *CAMmary* benchmark study [43], utilizing the identical ESA Collision Avoidance Challenge dataset introduced previously. The performance of the algorithm is systematically evaluated across a comprehensive matrix of safety constraints, assessing both deterministic physical clearance and probabilistic risk thresholds. The specific operational parameters and target values evaluated in this study are defined as follows:

- **Maximum Control Acceleration:** A constant maximum available thrust acceleration of $a_{max} = 0.01$ m/s^2 is assumed for all simulated low-thrust maneuvers.
- **Target Miss Distances (MD):** The deterministic spatial safety constraints, $\bar{d}$, are evaluated at discrete thresholds of 0.5 km, 1.5 km, 2.5 km, and 3.5 km.
- **Target Probabilities of Collision (PoC):** The statistical safety constraints are evaluated at maximum acceptable risk limits of 10^{-4}, 10^{-5}, 10^{-6}, and 10^{-7}.

A critical variable in any CAM design is the maneuver lead time, defined by the initial maneuver epoch t_0, which dictates how far in advance of the Time of Closest Approach (TCA) the optimization begins. As noted in the benchmark study regarding their analytical methodologies: "Given that Bomb15 and DeVit24 need a specific firing time, a total of 10 possible manoeuvring points are considered, and the cheapest manoeuvre is selected" [43]. Because the proposed DC method relies on mapping state sensitivities via the State Transition Matrix from a fixed initial state, it identically requires a predefined maneuver start time to initiate the backpropagation and root-finding process. Consequently, to maintain direct methodological comparability with the baseline algorithms, the maneuver start times in this study are discretized into the exact same 10-point temporal grid utilized by [43]. These maneuver lead times represent backpropagation times with respect to TCA ranging from 0.1 to 1.0 orbital periods, evaluated in regular increments of 0.1 orbits. For

[1] `https://github.com/arma1978/conjunction`

each individual conjunction event within the dataset, the proposed differential correction algorithm is systematically executed across all 10 temporal grid points. The feasible solution yielding the absolute minimum propellant consumption (Δv) is subsequently selected as the definitive optimal maneuver for that specific encounter.

4.2. Miss Distance Constraints Results

This section evaluates the performance of the Differential Correction (DC) methodology when tasked with enforcing a deterministic physical separation, or Miss Distance (MD), between the primary and secondary objects at the Time of Closest Approach (TCA). The results demonstrate remarkable accuracy and convergence rates across the dataset. It is worth noting that the number of actively evaluated samples varies depending on the targeted miss distance threshold (e.g., 786 samples for 0.5 km, 1846 for 1.5 km, and all 2170 for the 2.5 km and 3.5 km cases). This discrepancy arises because, for smaller target thresholds, many nominal conjunctions already satisfy the safety constraint without the need for a collision avoidance maneuver.

4.2.1. Convergence and Accuracy

The DC method achieves a flawless 100% convergence rate across all tested samples, all miss distance targets, and all maneuver lead times (from 0.1 to 1.0 orbital periods before TCA). Furthermore, the algorithm demonstrates exceptional precision. A strict miss distance error tolerance of 10^{-4} km (0.1 m) is imposed. The method successfully satisfies this tolerance in 100% of the cases for maneuvers initiated between 0.2 and 1.0 orbital periods prior to TCA across all distance categories. The only exceptions occur in the most time-constrained scenarios, specifically when maneuvers are initiated just 0.1 orbital periods before TCA:

- **2.5 km Target:** 9 out of 2170 cases (99.6% accuracy) slightly exceed the strict tolerance, with a maximum registered error of 1.35 m.
- **3.5 km Target:** 13 out of 2170 cases (99.4% accuracy) exceed the tolerance, with a maximum registered error of 1.33 m.

These minor deviations occur because finding a highly precise trajectory becomes inherently more difficult when operating close to TCA, as the low-thrust propulsion unit has severely limited time to integrate the required separation distance. However, it must be emphasized that these maximum tolerance violations correspond to positional errors of less than 1.4 meters. In the context of an operational maneuver targeting a multi-kilometer

miss distance, a sub-two-meter error is entirely negligible and does not invalidate the operational safety or the mathematical robustness of the proposed method.

4.2.2. Propellant Consumption: Minimum Required ΔV Analysis

To assess the fuel optimality of the proposed DC method, violin plots are generated to visualize the distribution of the *minimum required velocity change (ΔV) across all samples for each miss distance threshold*. The analysis of the propellant consumption distributions and computational efficiency yields the following insights, additionally supported by the mean values presented in Table 4.1:

- **0.5 km Miss Distance**, Figure 4.1: The DC method registers a maximum ΔV of approximately 0.13 m/s, with the mode (most common value) remaining well below 0.02 m/s. In terms of average expenditure, the DC solver achieves a mean ΔV of 45.32 mm/s. At this threshold, the DC method clearly outperforms DeVit23 (mean 50.98 mm/s) and performs similarly to Pav24b in terms of peak ΔV, while remaining highly competitive with the analytical methods. The algorithm demonstrates excellent computational speed, requiring a mean computation time of just 0.119 s per solution.
- **1.5 km Miss Distance**, Figure 4.2: The maximum registered ΔV increases to approximately 0.41 m/s, with the mode residing around 0.07 m/s. The average propellant consumption scales to 146.4 mm/s. The proposed algorithm again exhibits superior fuel optimality compared to DeVit23 (mean 164.4 mm/s) and yields results virtually identical to Pav24b in terms of peak ΔV. The mean computation time remains highly efficient at 0.148 s.
- **2.5 km Miss Distance**, Figure 4.3: For this larger separation, the maximum ΔV is slightly below 0.7 m/s, with the majority of samples clustering around a mode of 0.12 m/s. The mean required ΔV across the dataset is 271.7 mm/s. Notably, in this regime, the DC method outperforms both DeVit23 and Pav24b regarding the maximum required ΔV, while easily maintaining a much lower average cost than DeVit23 (309.5 mm/s). Finding these larger separation trajectories requires an average of 0.164 s per maneuver.
- **3.5 km Miss Distance**, Figure 4.4: In the most demanding scenario, the peak ΔV approaches 0.98 m/s, while the most common value remains below 0.2 m/s. The overall mean propellant cost reaches 420.2 mm/s. Similar to the 2.5 km case,

the DC method successfully places itself as a superior solver, outperforming both DeVit23 and Pav24b in terms of maximum fuel expenditure. It also maintains an average fuel consumption substantially lower than DeVit23 (483.2 mm/s) and close to the remaining benchmark models. It also demonstrates a rapid and robust mean convergence time of 0.189 s.

Method	MD = 0.5 km	MD = 1.5 km	MD = 2.5 km	MD = 3.5 km
Bomb15 [8]	41.26 mm/s	134.4 mm/s	252.9 mm/s	395.1 mm/s
DeVit24 [20]	41.34 mm/s	133.9 mm/s	252.4 mm/s	394.6 mm/s
DeVit23 [19]	50.98 mm/s	164.4 mm/s	309.5 mm/s	483.2 mm/s
Pav24b [44]	41.89 mm/s	137.8 mm/s	260.7 mm/s	411.8 mm/s
Arm21 [6]	41.37 mm/s	133.9 mm/s	251.8 mm/s	395.5 mm/s
Pav24a [41]	41.38 mm/s	133.8 mm/s	251.8 mm/s	393.6 mm/s
DC	45.32 mm/s	146.4 mm/s	271.7 mm/s	420.2 mm/s

Table 4.1: Mean value of the ΔV distribution for each considered miss distance threshold

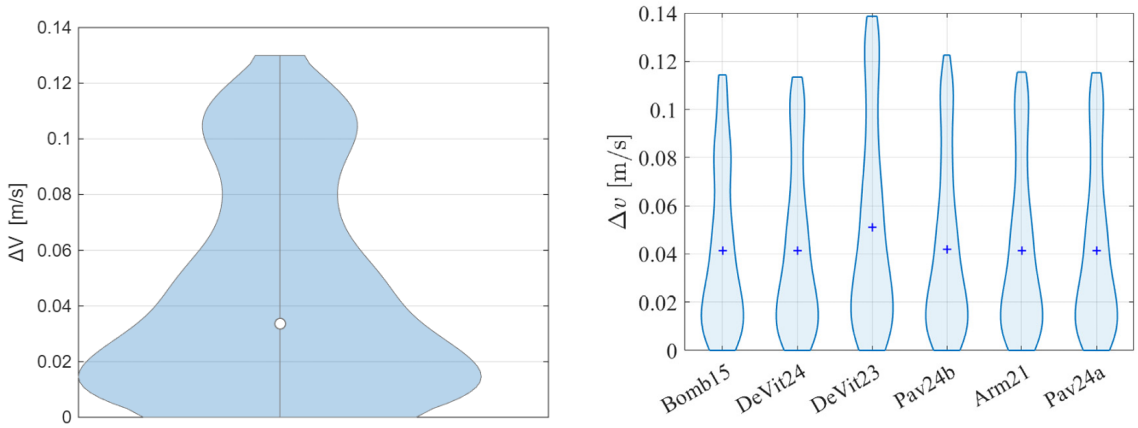

(a) DC method.

(b) CAMmary methods.

Figure 4.1: Violin plots illustrating the distribution of the minimum Δv required for target miss distances of 0.5 km

Overall, the evaluation of the deterministic miss distance constraints demonstrates the incredible robustness and consistency of the DC method. The algorithm seamlessly matches the performance of the best contemporary solvers and even outperforms them in scenarios requiring large separation distances. Finally, as theoretically expected, both the maximum and mean Δv values increase monotonically with the imposed miss distance, reflecting

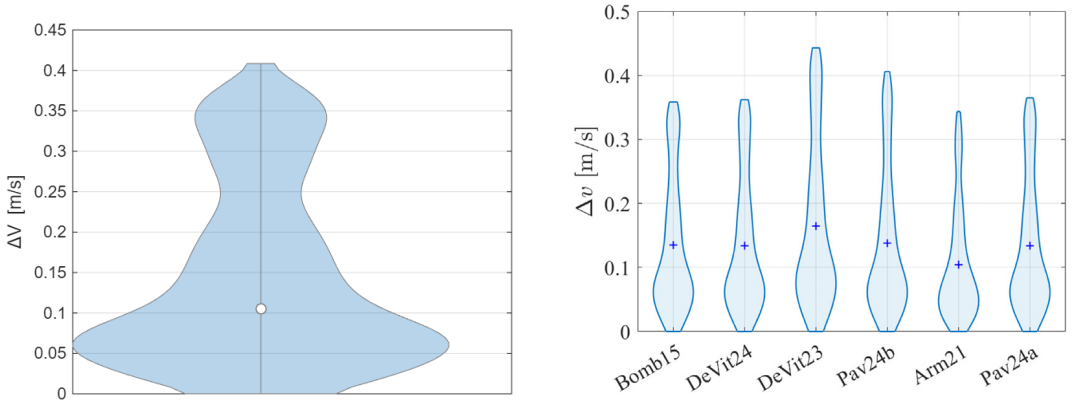

(a) DC method. (b) CAMmary methods.

Figure 4.2: Violin plots illustrating the distribution of the minimum Δv required for target miss distances of 1.5 km

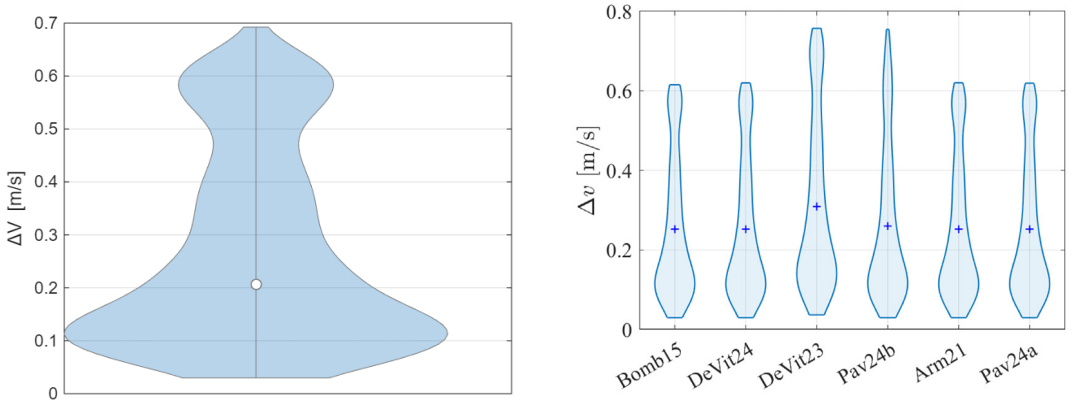

(a) DC method. (b) CAMmary methods.

Figure 4.3: Violin plots illustrating the distribution of the minimum Δv required for target miss distances of 2.5 km

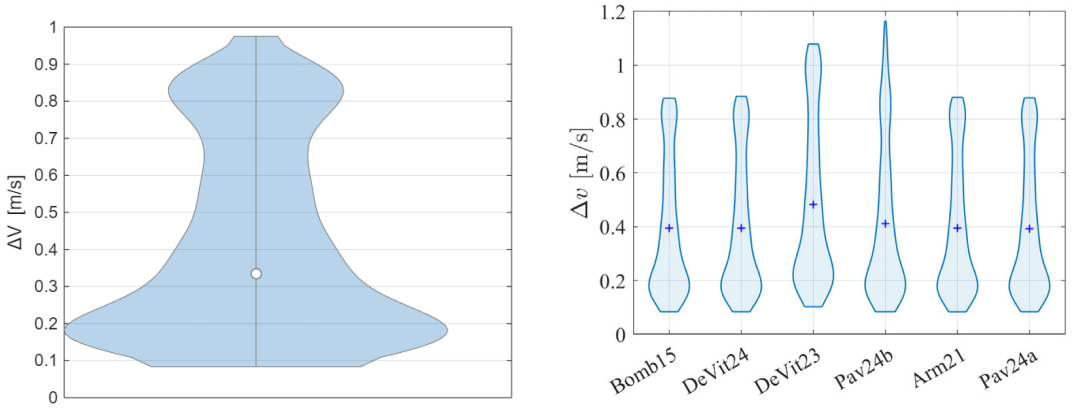

(a) DC method. (b) CAMmary methods.

Figure 4.4: Violin plots illustrating the distribution of the minimum Δv required for target miss distances of 3.5 km

the fundamental orbital mechanics of the problem: a larger required physical separation at TCA demands a greater expenditure of kinetic energy.

4.2.3. Optimal Maneuver Lead Time Analysis

An additional parameter of interest evaluated during the numerical campaign is the preferred firing time (or maneuver lead time) selected by the optimization process for each miss distance constraint. In orbital mechanics, the efficiency of a low-thrust maneuver is heavily dependent on the true or mean anomaly at which the thrust is applied relative to the Time of Closest Approach (TCA). Let $\Delta\theta$ represent the angular orbital separation (in degrees) between the maneuver initiation epoch t_0 and the TCA (t_{TCA}). Assuming a near-circular orbit with mean motion n, this can be approximated as:

$$\Delta\theta \approx n(t_{TCA} - t_0) \times \frac{180^\circ}{\pi}$$

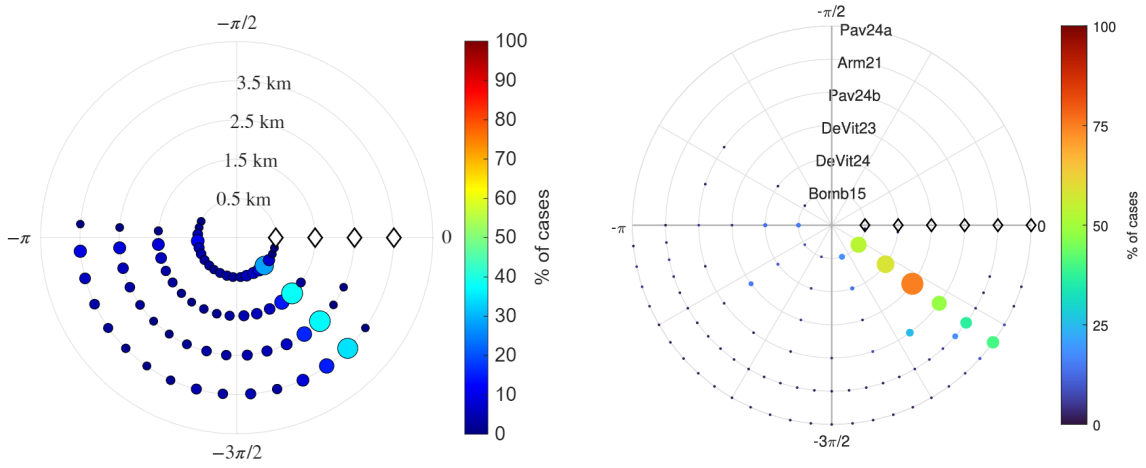

(a) DC method. (b) CAMmary methods.

Figure 4.5: Preferred firing times by each considered method: TCA is represented by the rhomboid at 0; firing close to 0 counter-clockwise means that the manoeuvre happens close to one revolution before TCA

Across the dataset, the optimization algorithm consistently favors initiating the maneuver at a specific phase to maximize the geometric displacement on the B-plane per unit of propellant mass. The results indicate that the most common optimal firing time occurs at $\Delta\theta \approx 315^\circ$ before TCA (which corresponds to approximately 0.875 orbital periods prior to conjunction). This optimal phasing aligns remarkably well with the theoretical and nu-

merical expectations established in the literature and it is clearly visualized in Figure 4.5. Specifically, the preferred value of 315° is highly consistent with the benchmark optimal firing time of 324° before TCA obtained by the analytical and numerical methodologies evaluated in the *CAMmary* review [43]. This indicates that the proposed DC method reliably identifies the globally optimal dynamic phasing for separation maneuvers.

4.2.4. Cumulative Distribution of Constraint Violations

To further quantify the precision of the proposed Differential Correction (DC) algorithm, the Cumulative Distribution Function (CDF) of the safety constraint violations is evaluated. In this context, a constraint violation, denoted as ε_d, is defined as the shortfall between the target deterministic miss distance $\bar{d}$ and the achieved miss distance d. Mathematically, this is expressed as:

$$\varepsilon_d = \max(0, \bar{d} - d)$$

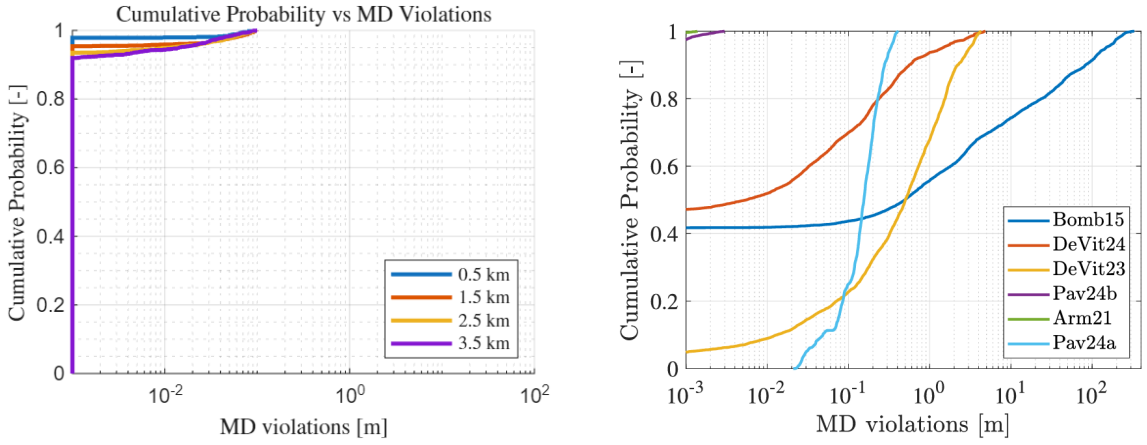

(a) DC method. (b) CAMmary methods.

Figure 4.6: Cumulative Distribution Function (CDF) of the miss distance constraint violations across all tested samples, comparing the proposed DC method to the *CAMmary* benchmarks

Consistent with the rigorous benchmarking methodology established in [43], analyzing the CDF of these violations provides a robust statistical measure of the solver's reliability. A highly capable algorithm is characterized by a steep CDF curve that rapidly converges to 1.0 (or 100%) at a near-zero error threshold, indicating that the vast majority of optimized maneuvers successfully satisfy or safely exceed the required spatial constraints. The analysis is visualized in Figure 4.6 and reveals that the DC method exhibits extraordinary precision. When evaluating the selected fuel-optimal maneuver profiles across all samples,

100% of the cases successfully restrict the miss distance error to strictly less than 0.1 m (i.e., $\epsilon < 10^{-4}$ km). When these CDF metrics are superimposed against the baseline solvers evaluated in [43], the proposed DC method firmly places itself as the second most accurate method overall. It provides an optimal balance between computational convergence and strict constraint adherence, ensuring that operational safety thresholds are met without the excessive constraint relaxation often required by competing low-thrust optimization algorithms.

4.2.5. Visualization of the First Sample: ID 1

To provide a visual representation of the algorithm's performance, the first sample of the dataset (ID 1) is analyzed in detail. For this case, the maneuver lead time (backpropagation time) is set to *half an orbital period before the Time of Closest Approach (TCA)*, with a target deterministic miss distance of 3.5 km.

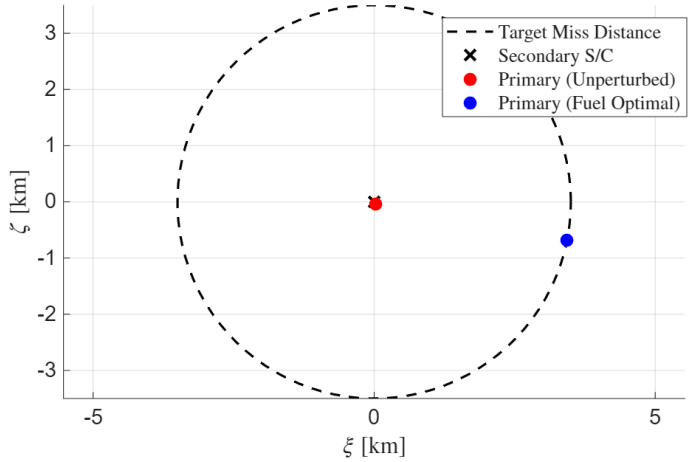

Figure 4.7: Visualization of the B-plane comparing the satellite positions before and after the maneuver

Figure 4.7 illustrates the B-plane positions of the primary and secondary satellites both before and after the maneuver. Prior to the maneuver, the two satellites are dangerously close, with an initial miss distance of approximately 43 m. Following the maneuver, the primary satellite is successfully displaced to exactly match the imposed safety threshold. Furthermore, it is important to note that the target miss distance region on the B-plane forms a perfect circle. As is demonstrated later, this geometric characteristic does not hold when imposing a squared Mahalanobis distance constraint.

Figure 4.8 presents the evolution of the costate vector components over time. Focusing on the detailed view of the burning window, it is evident that the initial guess for the optimal thrusting times is highly accurate. The optimization algorithm converges to a

window situated very close to the initial estimate, primarily shifting the ignition to a slightly later epoch without significantly altering the total burn duration. While this particular case exhibits only a single intersection of the costate magnitude $||\lambda(t)||$ with the switching threshold λ_{th}, the profile of the costate vector can vary significantly across different conjunction scenarios. Indeed, even for the same sample ID, altering the maneuver lead time can result in multiple threshold intersections located in different regions of the backpropagation envelope.

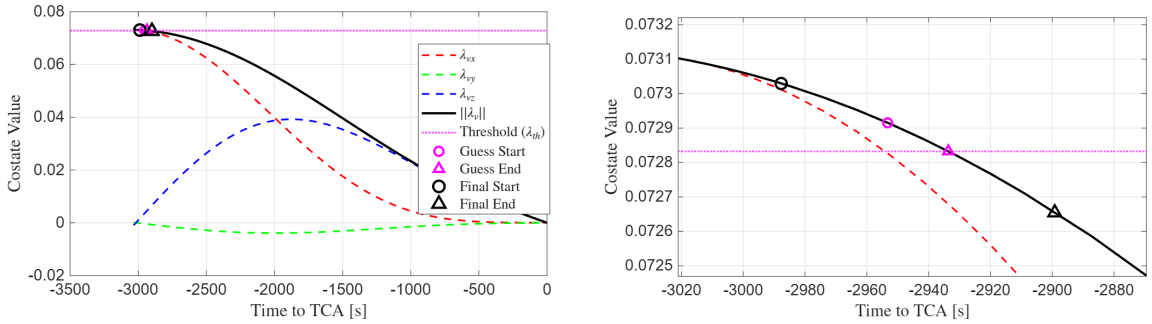

(a) Full evolution. (b) Burning window detail.

Figure 4.8: Evolution of the costate vector components over time for sample ID 1 (backpropagation time of 0.5 orbits before TCA), including a detailed view of the burning window

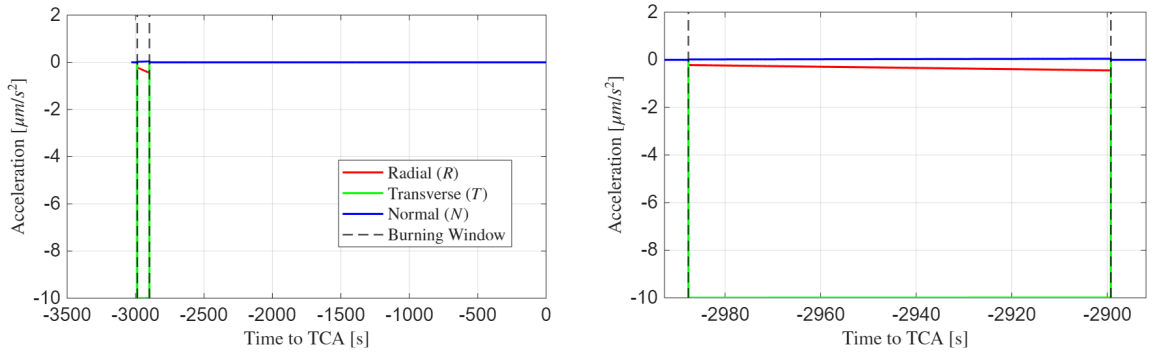

(a) Full acceleration profile. (b) Burning window detail.

Figure 4.9: Acceleration components over time for sample ID 1 (backpropagation time of 0.5 orbits before TCA), including a detailed view of the burning window

Finally, Figure 4.9 depicts the commanded acceleration over time. The control profile clearly exhibits a perfect bang-bang structure, characterized by a single, finite burning window. Notably, the thrust is primarily directed along the tangential axis of the Radial-Transverse-Normal (RTN) frame, accompanied by a minor radial component. This behavior aligns with theoretical expectations and is consistent with the optimal control profiles

derived by other methodologies in the literature.

4.3. Squared Mahalanobis Distance Constraints Results

This section evaluates the performance of the Differential Correction (DC) methodology when tasked with enforcing a statistical safety threshold—specifically, a maximum allowable Probability of Collision (PoC) at the Time of Closest Approach (TCA). To achieve this efficiently, the algorithm exploits the fundamental mathematical duality between probabilistic risk and geometric separation on the B-plane. Imposing an upper bound on the PoC is functionally identical to enforcing a minimum limit on the Squared Mahalanobis Distance (d_M^2).

While the PoC defines the statistical likelihood of an intersection, the d_M^2 acts as its deterministic, geometric counterpart. It translates a complex probability integral into a tangible, dynamically scaled exclusion ellipse defined by the encounter's combined covariance. By utilizing d_M^2 as a strictly monotonic surrogate function, the optimizer is able to natively and robustly satisfy the PoC requirement through a well-defined geometric boundary condition.

The results demonstrate remarkable accuracy and convergence rates across the dataset. Similar to the deterministic cases, the number of actively evaluated samples varies depending on the targeted PoC threshold (e.g., 1230 samples for 10^{-4}, 1832 for 10^{-5}, 2103 for 10^{-6}, and 2133 for the 10^{-7} case). This discrepancy reflects reality: nominal conjunctions that already fall safely below the prescribed statistical risk threshold do not trigger a collision avoidance maneuver.

4.3.1. Convergence and Accuracy

The DC method achieves a flawless 100% convergence rate across all tested samples, all targeted PoC thresholds, and all maneuver lead times (ranging from 0.1 to 1.0 orbital periods before TCA). Furthermore, the algorithm demonstrates good precision in satisfying the statistical constraints, though underperforming with respect to the Miss Distance constraint cases. Unlike the deterministic miss distance constraints, much more outliers are observed among all possible backpropagation times. To ensure a realistic operational assessment, extreme propellant consumption outliers, in terms of accuracy and ΔV cost are filtered out prior to evaluating the mean statistics.

Certain conjunction geometries naturally dictate highly inefficient maneuvers (requiring

up to $\Delta V \approx 3$ m/s). This is not a failure of the algorithm, but a consequence of the unfavorable orientation of the covariance ellipses relative to the available thrust authority, forcing the differential corrector to find an accurate but prohibitively expensive solution. Filtering these unfeasible geometric outliers guarantees that the reported metrics accurately reflect the solver's capabilities for standard operational scenarios.

4.3.2. Propellant Consumption: Minimum Required ΔV Analysis

To assess the fuel optimality of the proposed DC method under probabilistic constraints, violin plots are generated to visualize the distribution of the minimum required velocity change (ΔV) across the filtered samples for each PoC threshold.

The analysis of the propellant consumption distributions and computational efficiency yields the following insights, additinally supported by the mean values presented in Table 4.2:

- **PoC 10^{-4} Threshold**, Figure 4.10: For this baseline risk level, outliers requiring $\Delta V \geq 0.08$ m/s are removed, retaining 97.9% of the samples (1204/1230). The DC method registers a mean ΔV of 14.11 mm/s (after filtering). The algorithm demonstrates excellent computational speed, requiring a mean computation time of 0.203 s per solution.
- **PoC 10^{-5} Threshold**, Figure 4.11: For this stricter threshold, outliers requiring $\Delta V \geq 0.16$ m/s are filtered out, retaining 96.3% of the samples (1764/1832). The average propellant consumption scales to 25.07 mm/s (after filtering). The mean computation time increases slightly to 0.237 s, reflecting the tighter integration tolerances required to resolve smaller target probabilities.
- **PoC 10^{-6} Threshold**, Figure 4.12: At this high-safety threshold, maintaining the 0.25 m/s filter retains 95.5% of the samples (2009/2103). The mean required ΔV across the dataset grows to 35.19 mm/s (after filtering). Extracting these trajectories requires an average computation time of 0.220 s per maneuver.
- **PoC 10^{-7} Threshold**, Figure 4.13: In the most demanding statistical scenario, outliers requiring $\Delta V \geq 0.35$ m/s are filtered out, retaining 95.8% of the samples (2043/2133), the overall mean propellant cost reaches 46.29 mm/s (after filtering). Despite the stringent requirements, the solver continues to demonstrate rapid and robust convergence, achieving a mean computation time of 0.217 s.

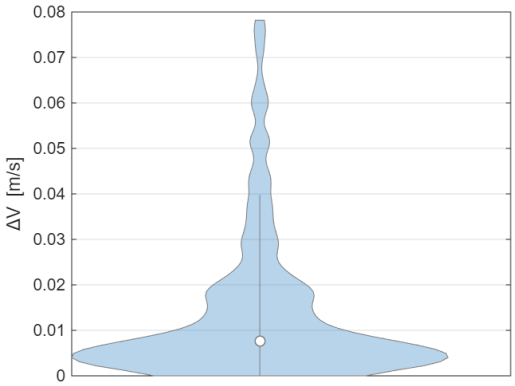

(a) DC method (after removing outliers).

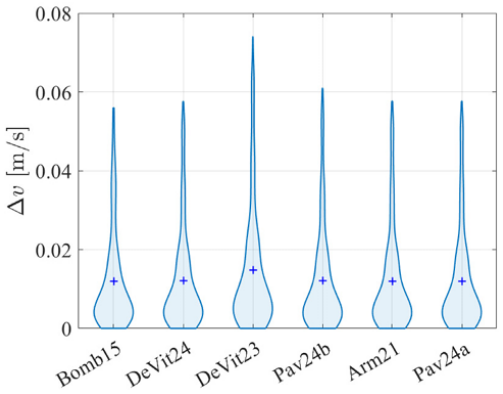

(b) CAMmary methods.

Figure 4.10: Violin plots illustrating the distribution of the minimum ΔV required for a target PoC of 10^{-4}

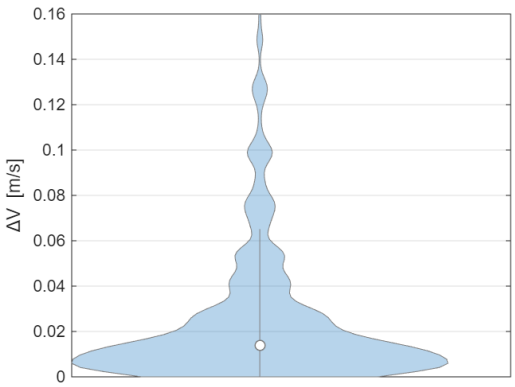

(a) DC method (after removing outliers).

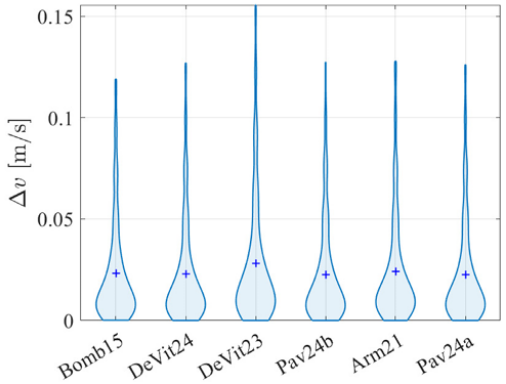

(b) CAMmary methods.

Figure 4.11: Violin plots illustrating the distribution of the minimum ΔV required for a target PoC of 10^{-5}

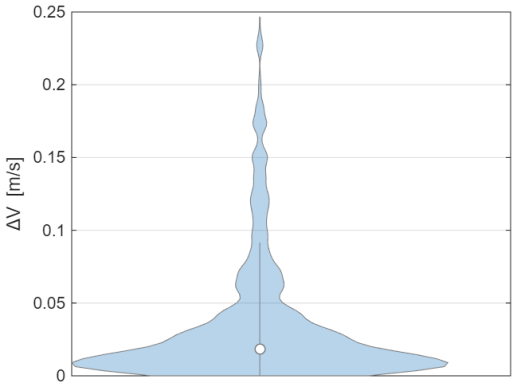

(a) DC method (after removing outliers).

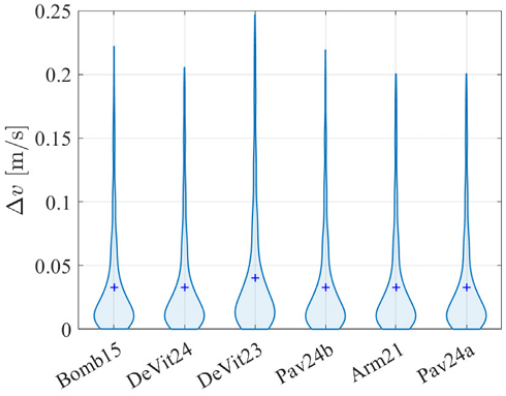

(b) CAMmary methods.

Figure 4.12: Violin plots illustrating the distribution of the minimum ΔV required for a target PoC of 10^{-6}

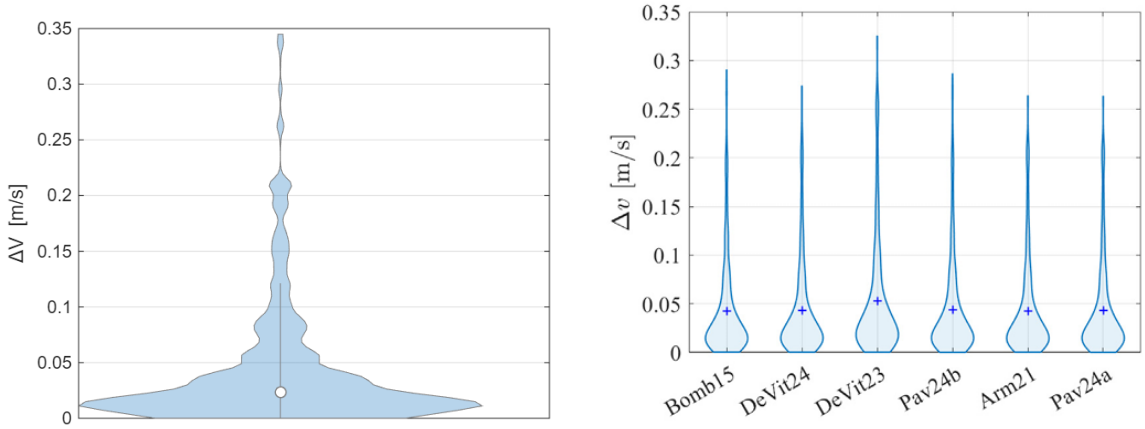

(a) DC method (after removing outliers). (b) CAMmary methods.

Figure 4.13: Violin plots illustrating the distribution of the minimum ΔV required for a target PoC of 10^{-7}

Method	**PoC** $= 10^{-4}$	**PoC** $= 10^{-5}$	**PoC** $= 10^{-6}$	**PoC** $= 10^{-7}$
Bomb15 [8]	11.98 mm/s	23.38 mm/s	32.73 mm/s	42.55 mm/s
DeVit24 [20]	12.03 mm/s	22.78 mm/s	33.00 mm/s	43.09 mm/s
DeVit23 [19]	14.80 mm/s	28.03 mm/s	40.50 mm/s	52.79 mm/s
Pav24b [44]	12.11 mm/s	22.70 mm/s	33.07 mm/s	43.80 mm/s
Arm21 [6]	11.96 mm/s	24.30 mm/s	32.84 mm/s	42.80 mm/s
Pav24a [41]	12.00 mm/s	22.75 mm/s	32.91 mm/s	42.90 mm/s
DC	14.11 mm/s	25.07 mm/s	35.19 mm/s	46.29 mm/s

Table 4.2: Mean value of the ΔV distribution for each considered PoC threshold

Overall, the evaluation of the Squared Mahalanobis Distance constraints validates the DC methodology's capability to natively handle complex, covariance-driven statistical bounds.
As theoretically expected, demanding a lower acceptable collision risk monotonically expands the d_M^2 boundary, which correspondingly increases the mean ΔV expenditure.

4.3.3. Optimal Maneuver Lead Time Analysis

An additional parameter of interest evaluated during the numerical campaign is the preferred firing time (or maneuver lead time) selected by the optimization process for each

statistical constraint.

Let $\Delta\theta$ represent the angular orbital separation (in degrees) between the maneuver initiation epoch t_0 and the TCA (t_{TCA}). Assuming a near-circular orbit with mean motion n, this can be approximated as:

$$\Delta\theta \approx n(t_{TCA} - t_0) \times \frac{180^\circ}{\pi}$$

Consistent with the deterministic miss distance results, the optimization algorithm favors initiating the maneuver at a specific phase to maximize the risk reduction per unit of propellant mass. Figure 4.14 shows the results and it can be easily noted that the most common optimal firing time remains structurally tied to the dynamics of the B-plane projection, occurring at $\Delta\theta \approx 180^\circ$ before TCA. This preferred phasing proves that the dynamic efficiency of the low-thrust arc is fundamentally preserved, regardless of whether the target boundary takes the form of an isotropic circle (miss distance) or a covariance-scaled ellipse (Mahalanobis distance).

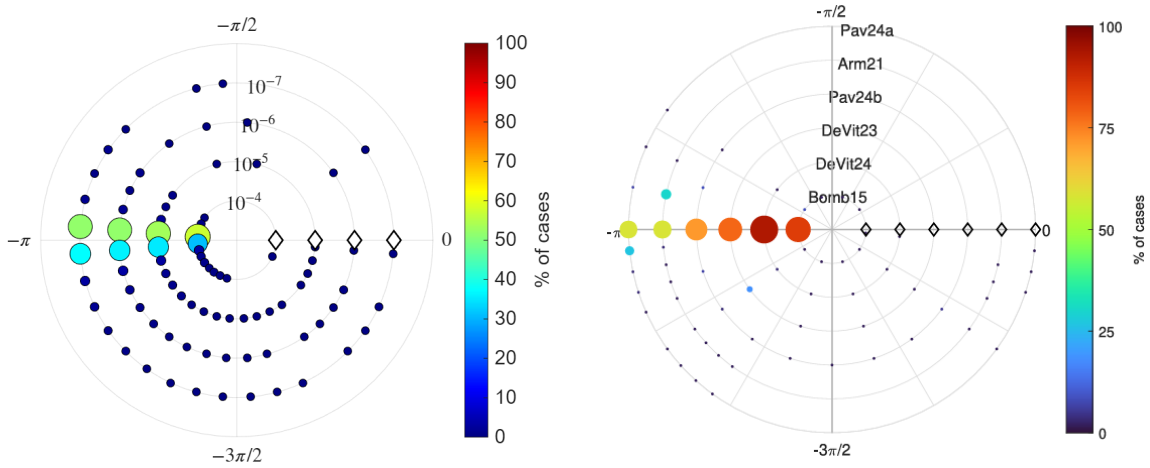

(a) DC method. (b) CAMmary methods.

Figure 4.14: Preferred firing times by each considered method for PoC targets: TCA is represented by the rhomboid at 0; firing close to 0 counter-clockwise means the maneuver happens close to one revolution before TCA

4.3.4. Cumulative Distribution of Constraint Violations

To quantify the precision of the proposed algorithm under probabilistic requirements, the Cumulative Distribution Function (CDF) of the safety constraint violations is evaluated. Because the algorithm targets a maximum threshold for the Probability of Collision, a

constraint violation is defined as the relative overshoot between the achieved PoC and the target $\overline{PoC}$. Mathematically:

$$\varepsilon_{PoC} = \max\left(0, \frac{PoC}{\overline{PoC}} - 1\right)$$

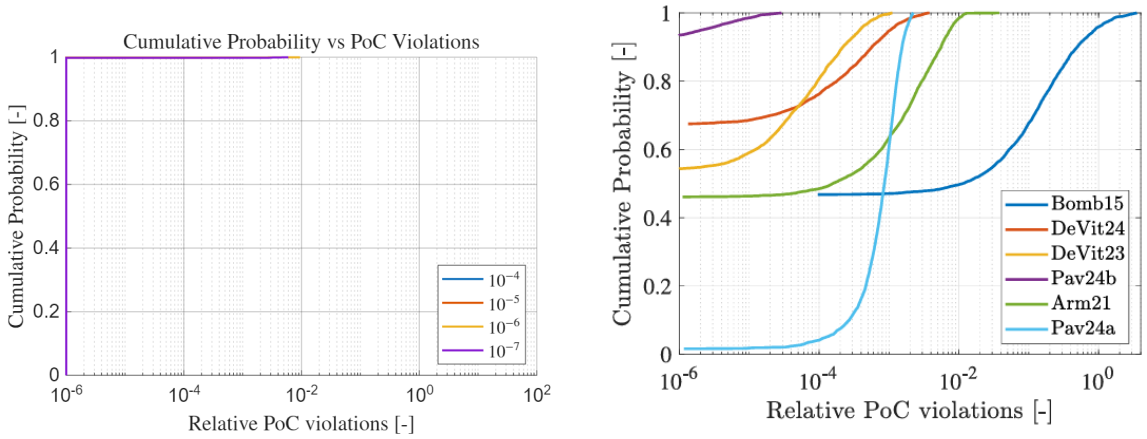

(a) DC method.

(b) CAMmary methods.

Figure 4.15: Cumulative Distribution Function (CDF) of the Probability of Collision constraint violations across all tested samples

A highly capable algorithm is characterized by a steep CDF curve that rapidly converges to 1.0 (or 100%) at a near-zero error threshold, indicating that the vast majority of optimized maneuvers successfully satisfy the required statistical constraints without penetrating the covariance-defined danger volume. The analysis, visualized in Figure 4.15, confirms the DC method's extraordinary precision in handling statistical boundaries, firmly achieving constraint adherence rates that easily parallel the benchmarks evaluated in [43].

4.3.5. Visualization of the First Sample: ID 1

To provide a visual representation of the algorithm's performance under statistical constraints, the first sample of the dataset (ID 1) is analyzed. For this case, the maneuver lead time (backpropagation time) is set to *half an orbital period before the Time of Closest Approach (TCA)*. Figure 4.18 maps the B-plane positions of the primary and secondary satellites both before and after the maneuver.

Here, the duality between PoC and d_M^2 is visually apparent: the target safety region is inherently probabilistic in definition, yet it materializes geometrically as an exclusion ellipse, scaled and oriented by the combined covariance matrix of the encounter. Following the maneuver, the primary satellite is successfully displaced to lie precisely on the border

of this d_M^2 ellipse, physically guaranteeing that the required maximum PoC is met without expending unnecessary propellant.

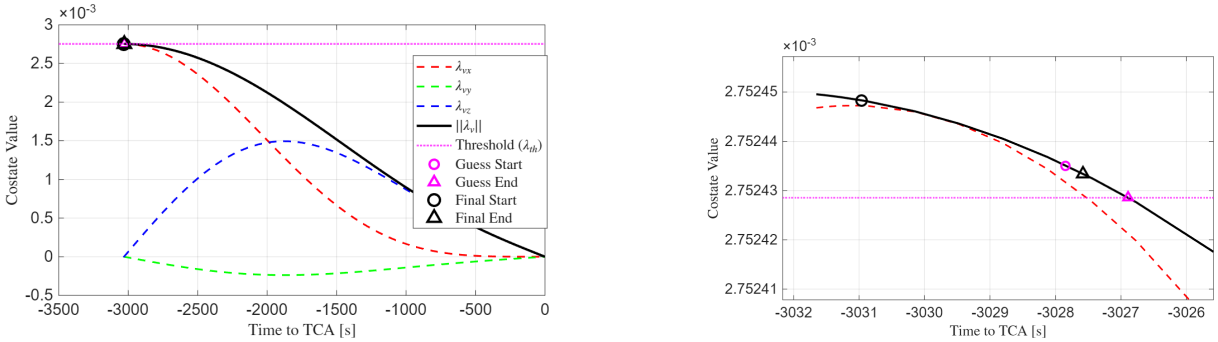

Figure 4.16: Evolution of the costate vector components over time for sample ID 1 under PoC constraints.

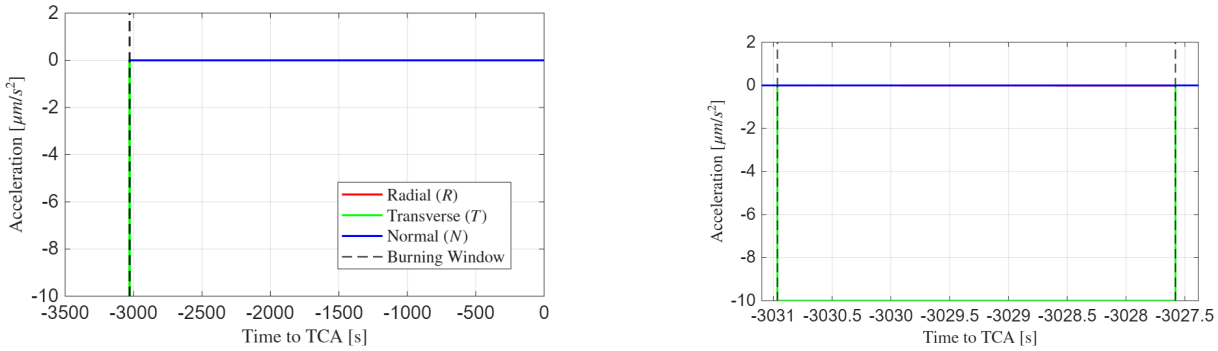

Figure 4.17: Acceleration components over time for sample ID 1 under PoC constraints.

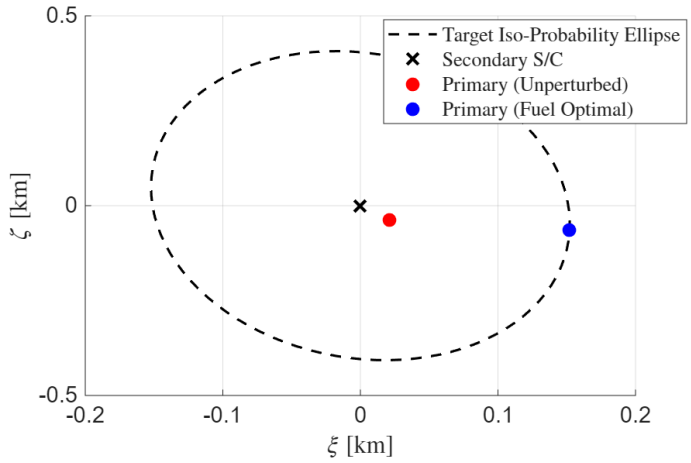

Figure 4.18: Visualization of the B-plane comparing satellite positions before and after the maneuver, showcasing the dual nature of the constraint: targeting a specific PoC physically manifests as an elliptical Squared Mahalanobis Distance boundary

Figure 4.16 presents the evolution of the costate vector components over time, while Figure 4.17 depicts the commanded acceleration. Even with the transition from an isotropic scalar distance to a covariance-mapped elliptical target boundary, the control profile preserves a distinct bang-bang structure with a finite burning window, primarily directed along the tangential axis.

4.4. Effects of Drag and J_2 Perturbations on Method Accuracy

To rigorously stress-test the accuracy and stiffness of the proposed Differential Correction (DC) method, the baseline optimal control profiles are re-evaluated under a high-fidelity dynamic environment. Because the optimization algorithm natively assumes unperturbed two-body problem (2BP) dynamics to ensure computational efficiency, it is crucial to quantify the constraint degradation when the spacecraft is subjected to realistic environmental perturbations—specifically, the Earth's oblateness (J_2) and atmospheric drag. A diagram of the process discussed in this section is shown in Figure 4.19.

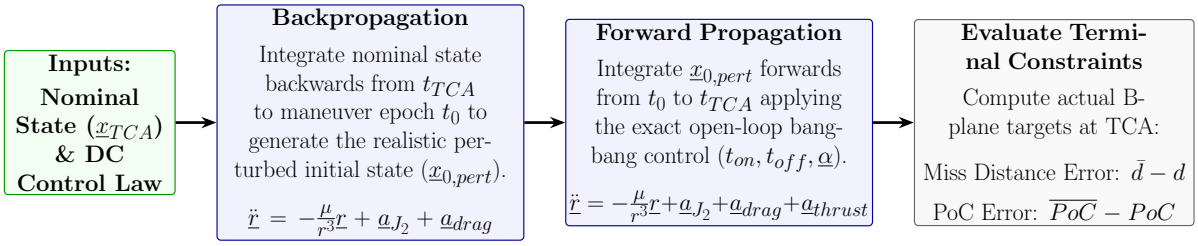

Figure 4.19: Perturbation effects validation diagram

The perturbed equations of motion governing the primary spacecraft are formulated by adjoining the disturbing accelerations to the nominal Keplerian dynamics:

$$\ddot{\mathbf{r}} = -\frac{\mu}{r^3}\mathbf{r} + \mathbf{a}_{J_2} + \mathbf{a}_{drag} + \mathbf{a}_{thrust}$$

The acceleration due to the Earth's J_2 oblateness is modeled analytically [15]:

$$\mathbf{a}_{J_2} = \frac{3}{2}\frac{J_2 \mu R_E^2}{r^4}\begin{bmatrix} \frac{x}{r}\left(5\frac{z^2}{r^2} - 1\right) \\ \frac{y}{r}\left(5\frac{z^2}{r^2} - 1\right) \\ \frac{z}{r}\left(5\frac{z^2}{r^2} - 3\right) \end{bmatrix}$$

where $R_E = 6371.0$ km represents the mean equatorial radius of the Earth, and $J_2 = 1.0826 \times 10^{-3}$ is the second zonal harmonic coefficient.

Simultaneously, the deceleration caused by atmospheric drag is evaluated using the standard aerodynamic drag model [15]:

$$\mathbf{a}_{drag} = -\frac{1}{2}\rho\frac{s_A}{m}C_D v_{rel}\mathbf{v}_{rel}$$

where ρ is the local atmospheric density computed using an exponential atmospheric model [60], C_D is the drag coefficient, s_A is the cross-sectional area, m is the spacecraft mass, and $\mathbf{v}_{rel}$ is the velocity vector of the spacecraft relative to the rotating atmosphere, with v_{rel} being its magnitude. This relative velocity is explicitly defined as:

$$\mathbf{v}_{rel} = \mathbf{v} - \omega_E \times \mathbf{r}$$

where $\mathbf{v}$ and $\mathbf{r}$ are the inertial velocity and position vectors of the spacecraft, and $\omega_E = [0, 0, \omega_E]^T$ is the Earth's angular velocity vector, with $\omega_E \approx 7.2921 \times 10^{-5}$ rad/s.

The validation procedure is structured as follows. First, the nominal state at the Time of Closest Approach (TCA) is backpropagated to the maneuver initiation epoch (t_0) using these perturbed dynamics to generate a realistic perturbed initial state, $\mathbf{x}_{0,pert}$. Subsequently, the controlled trajectory is forward-propagated from t_0 to TCA under the same perturbed dynamics, applying the exact same open-loop bang-bang control law (ignition times, cutoff times, and thrust directions) generated by the unperturbed DC method. Finally, the terminal miss distance and Probability of Collision (PoC) constraints are evaluated.

To emphasize the impact of the perturbations and create an extremely conservative testing environment, several worst-case assumptions are incorporated:

- **Atmospheric Density:** The local atmospheric density (ρ) is computed using an exponential atmospheric model derived from Wertz [60].
- **Aerodynamic Properties:** The active spacecraft is modeled with a small mass of $m = 50$ kg and a worst-case drag coefficient of $C_D = 2.6$ to heavily magnify the deceleration effects.
- **Cross-Sectional Area:** Because the ESA dataset only provides the combined cross-sectional area of the primary and secondary objects (s_A), this total combined value is conservatively assigned entirely to the primary evading spacecraft. This deliberate overestimation artificially amplifies the drag perturbations, providing a rigorous worst-case proof of the method's accuracy.

With this conservative perturbed dynamic environment established, the resulting devia-

tions in the terminal optimal trajectories are systematically quantified below.

4.4.1. Impact on Miss Distance Constraints

The performance of the open-loop DC control law under perturbed dynamics for deterministic target miss distances is summarized below:

- **0.5 km Target:** Out of all tested encounters, 576 cases (73.3%) resulted in an error exceeding the strict 10^{-4} km tolerance limit. However, the maximum registered error with respect to the target miss distance was only 1.65 m, while the overall mean error stood at a mere 0.42 m.
- **1.5 km Target:** The analysis identified 1760 cases (95.3%) technically violating the error tolerance (10^{-4} km). The maximum observed error was 5.01 m, and the mean error across all samples was 1.66 m.
- **2.5 km Target:** Tolerance violations (10^{-4} km) occurred in 2121 cases (97.7%). For these larger separations, the maximum error reached 8.34 m, accompanied by a mean error of 3.28 m.
- **3.5 km Target:** Almost all samples (2136 cases, representing 98.4%) exceeded the error tolerance (10^{-4} km). Nevertheless, the peak error remained extremely low at 11.76 m, with a mean error of 5.17 m.

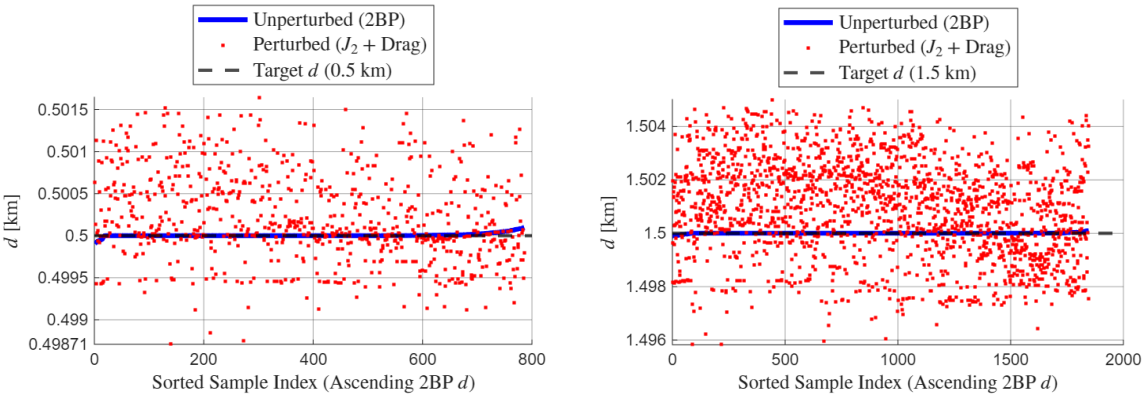

(a) Target Miss Distance: 0.5 km. (b) Target Miss Distance: 1.5 km.

Figure 4.20: Comparison of perturbed and unperturbed miss distance values across all dataset samples for target miss distances of 0.5 km and 1.5 km. The x-axis represents the sorted sample index, and the y-axis represents the achieved miss distance.

Although a significant percentage of the optimized trajectories now technically violate the previously imposed strict computational tolerance, the absolute magnitudes of these

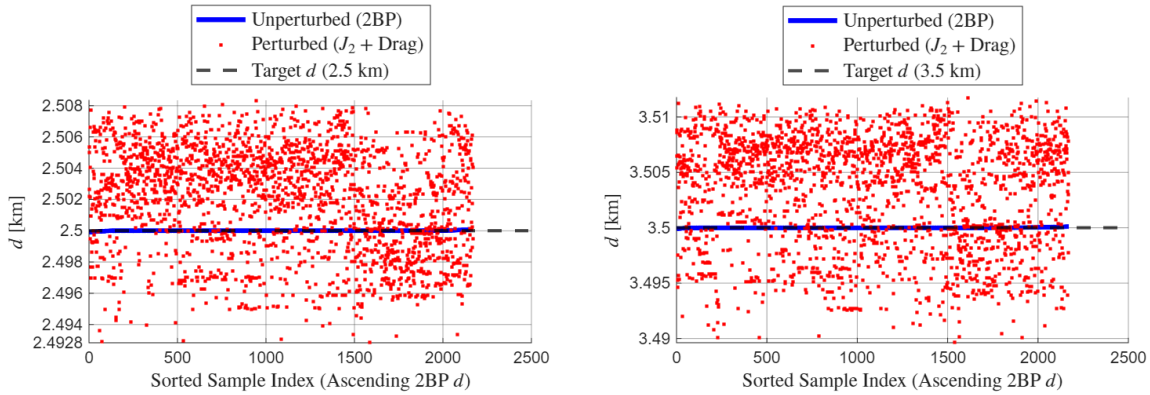

(a) Target Miss Distance: 2.5 km. (b) Target Miss Distance: 3.5 km.

Figure 4.21: Comparison of perturbed and unperturbed miss distance values across all dataset samples for target miss distances of 2.5 km and 3.5 km. The x-axis represents the sorted sample index, and the y-axis represents the achieved miss distance.

deviations remain remarkably small. The individual miss distance variations for all tested samples, sorted in ascending order based on their nominal 2BP values for clear visualization, are displayed in Figures 4.20 and 4.21. In the context of orbital maneuvers targeting multi-kilometer separations, meter-level deviations are operationally negligible. This confirms that evaluating the control law under unperturbed 2BP assumptions is a highly valid and computationally efficient strategy, as the spatial divergence caused by unmodeled perturbations over a maximum maneuver lead time of one orbit is practically insignificant.

4.4.2. Impact on Probability of Collision Constraints

The evaluation is similarly extended to the probabilistic PoC boundaries to test the stability of the covariance-scaled constraints. The individual PoC variations for all tested samples are illustrated in Figures 4.22 and 4.23.

The relative error tolerance for these statistical targets is set at a strict 5%.

- 10^{-4} **and** 10^{-5} **PoC Targets:** Zero cases violated the predefined 5% relative error tolerance, demonstrating perfect constraint adherence despite the perturbations.
- 10^{-6} **PoC Target:** Only 1 single case exceeded the 5% relative error margin.
- 10^{-7} **PoC Target:** Merely 2 cases recorded an error slightly above the 5% relative threshold.

For all assessed cases, the introduction of perturbed dynamics yields minimal deviations

relative to both the target PoC and the nominal unperturbed DC solutions. The errors remain strictly bounded within the set relative tolerance threshold, with the exception of only 3 borderline cases out of thousands tested. This highlights the extreme robustness of the algorithm when targeting complex statistical boundaries.

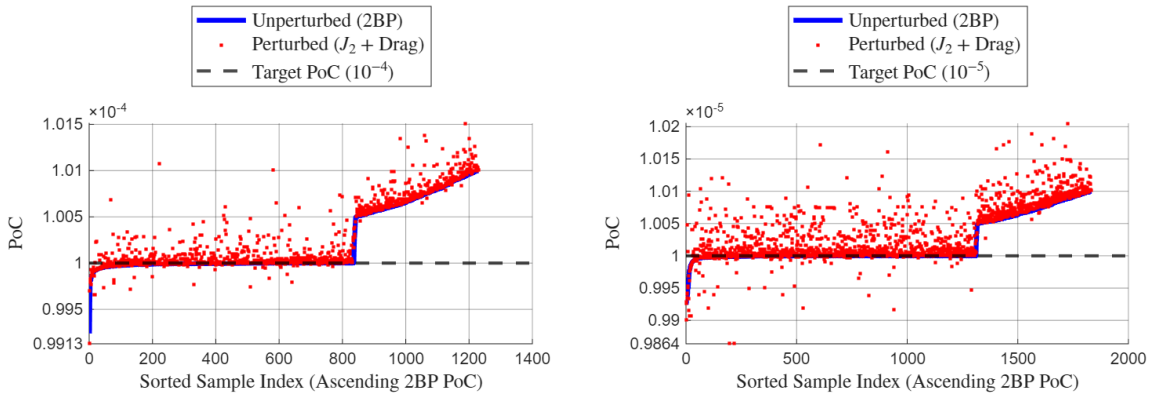

(a) Target PoC: 10^{-4}. (b) Target PoC: 10^{-5}.

Figure 4.22: Comparison of perturbed and unperturbed Probability of Collision (PoC) values across all dataset samples for target PoC thresholds of 10^{-4} and 10^{-5}. The x-axis represents the sorted sample index, and the y-axis represents the achieved PoC.

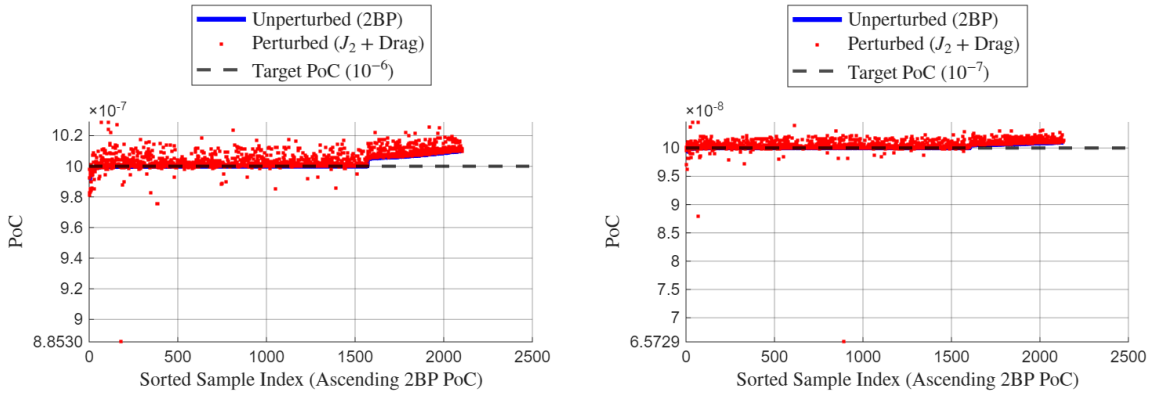

(a) Target PoC: 10^{-6}. (b) Target PoC: 10^{-7}.

Figure 4.23: Comparison of perturbed and unperturbed Probability of Collision (PoC) values across all dataset samples for target PoC thresholds of 10^{-6} and 10^{-7}. The x-axis represents the sorted sample index, and the y-axis represents the achieved PoC.

4.5. Summary of This Chapter

This chapter presents a comprehensive numerical validation of the proposed Differential Correction (DC) methodology for low-thrust Collision Avoidance Maneuvers. By evaluat-

ing the algorithm against established state-of-the-art models using the standardized ESA Collision Avoidance Challenge dataset, the analysis confirms the DC framework's exceptional robustness, accuracy, and computational efficiency. The solver successfully demonstrates its capability to natively enforce both physical and statistical safety boundaries with sub-second execution speeds, explicitly prioritizing the minimization of propellant consumption across a highly diverse catalog of short-term encounter scenarios.

1) **Deterministic Miss Distance Constraints:** The DC method achieves a flawless 100% convergence rate across all tested samples. It enforces strict physical separation thresholds (ranging from 0.5 km to 3.5 km) with remarkable precision, successfully maintaining positional errors below 0.1 meters in the vast majority of scenarios. Furthermore, the algorithm is even able to outperform some contemporary benchmark solvers in terms of fuel optimality, particularly at larger required separation distances, while maintaining an average computational time of less than 0.2 seconds per maneuver.

2) **Probabilistic Safety Constraints (Squared Mahalanobis Distance):** By exploiting the mathematical duality between the Probability of Collision (PoC) and the dynamically scaled d_M^2 boundary, the algorithm effectively targets complex, covariance-driven exclusion ellipses, achieving a 100% convergence rate across all tested statistical targets (10^{-4} to 10^{-7}). However, the method noticeably underperforms in this domain compared to the deterministic miss distance cases, generating a much higher number of extreme propellant consumption outliers. This underperformance is not a failure of the algorithm itself, but a direct consequence of the unfavorable orientation of the covariance ellipses relative to the available thrust authority in certain conjunction geometries, which forces the differential corrector to find accurate but prohibitively expensive trajectories. Once these unfeasible geometric outliers are filtered out to ensure a realistic operational assessment, the algorithm consistently yields highly competitive mean ΔV expenditures, successfully and robustly resolving complex probabilistic risk limits for standard operational scenarios.

3) **Optimal Maneuver Lead Times and Constraint Adherence:** The optimization process reliably identifies the globally optimal dynamic phasing for separation maneuvers. The results indicate a preferred firing time at approximately 315° before the Time of Closest Approach (TCA) for deterministic targets, and 180° for probabilistic boundaries. Additionally, the cumulative distribution function (CDF) of the constraint violations confirms the solver's supreme reliability, proving that it meets strict operational safety thresholds without requiring the excessive constraint relaxation often observed in standard low-thrust optimization algorithms.

4) **High-Fidelity Sample Visualization:** An in-depth analysis of test sample ID 1 visually confirms the method's operational precision. For a deterministic 3.5 km target *initiated half an orbit before conjunction*, the DC algorithm successfully displaces the primary satellite from a critical 43-meter initial miss distance to exactly match the required circular safety boundary. Similarly, under a strict 10^{-7} PoC target, the maneuver precisely positions the satellite exactly on the border of the scaled d_M^2 exclusion ellipse. In all visualized cases, the derived control profiles maintain a perfect, fuel-optimal bang-bang structure with highly accurate ignition and cutoff times and mainly tangential acceleration contribution, as noted by other solvers in the literature.

5) **Robustness Against Environmental Perturbations:** To rigorously validate the operational stiffness of the unperturbed 2BP baseline optimization, the derived open-loop control laws were re-evaluated under a high-fidelity dynamic environment incorporating J_2 Earth oblateness and an aggressively modeled, worst-case atmospheric drag. The results empirically demonstrate that although minor target boundary deviations naturally occur, they remain operationally negligible over the short-term maneuvering window (maximum one orbital period). For deterministic constraints, maximum physical deviations peaked at merely 11.76 meters (against a 3.5 km target), while probabilistic constraints maintained an exceptional adherence, with only 3 out of thousands of test cases slightly exceeding a 5% relative error threshold. This confirms that the unperturbed optimization strategy safely and accurately mitigates collision risks while successfully avoiding the computational burden of propagating complex perturbations.

5 | Conclusions and Future Work

This chapter summarizes the results of the present work and offers recommendations for future research in this field.

To address the critical need for autonomous trajectory planning in increasingly congested Low Earth Orbits, this thesis developed a highly mass-efficient Differential Correction (DC) methodology for low-thrust Collision Avoidance Maneuvers. By mapping initial state sensitivities directly to the encounter B-plane via the State Transition Matrix, the optimal control problem was successfully reformulated to yield a precise, two-switch bang-bang thrust profile. Specifically, the core research, mathematical formulations, and operational evaluations of this study achieved the following primary milestones:

Low-Thrust Satellite Collision Avoidance Modelling was formulated to establish the rigorous dynamical and statistical foundation of the encounter geometry. This model mapped the fully decoupled low-thrust two-body dynamics from the global Earth-Centered Inertial (ECI) frame into the local Radial-Transversal-Normal (RTN) frame and the encounter-specific B-plane. By back-propagating the nominal European Space Agency (ESA) challenge dataset, the severe geometric limitations of unperturbed trajectories were quantified. This robust mathematical mapping enabled the exact construction of the optimal control boundary conditions, allowing the solver to seamlessly transition between deterministic Euclidean Miss Distance limits (d) and covariance-scaled Squared Mahalanobis Distance limits (d_M^2) using Chan's equivalent cross-sectional area formulation.

Fuel Optimal Bang-Bang Control Using Differential Correction was developed to bypass the severe numerical instabilities associated with integrating discontinuous costates in traditional Two-Point Boundary Value Problems. This was achieved by freezing the thrust vector to the analytical energy-optimal direction—derived via a closed-form fourth-order polynomial root-finding process—and extracting initial switching bounds through a rapid bisection method on the velocity costate magnitude. A 42-dimensional differential system was then concurrently propagated to populate a 1×2 Jacobian matrix, leveraging the Moore-Penrose pseudoinverse to iteratively correct the engine ignition (t_{on}) and cutoff (t_{off}) times until the spatial B-plane constraints were strictly satisfied.

Performance Evaluation of the Low-Thrust Fuel Optimal Differential Correction was conducted to rigorously validate the operational efficiency, accuracy, and fuel optimality of the proposed methodology using 2,170 Conjunction Data Messages from the European Space Agency Collision Avoidance Challenge dataset, benchmarking the algorithm against 6 state-of-the-art analytical and numerical solvers. Through this extensive validation, the DC framework demonstrated exceptional robustness across all 2,170 samples and 10 temporal maneuver grid points.

For Deterministic Miss Distance (MD) constraints, the solver achieved a 100% convergence rate and satisfied the strict 0.1-meter error tolerance in 100% of cases for maneuvers initiated beyond 0.2 orbital periods prior to TCA. Under extreme time constraints (0.1 orbits), accuracy exceeded 99.4% with maximum positional errors strictly under 1.4 meters. Mean ΔV scaled efficiently from 45.32 mm/s (for a 0.5 km target) to 420.2 mm/s (for a 3.5 km target), maintained rapid execution with mean computation times ranging from 0.119 to 0.189 seconds, and autonomously identified maximum dynamic efficiency at approximately 315 degrees before TCA.

For Probabilistic Probability of Collision (PoC) constraints, the solver similarly achieved a 100% convergence rate, reliably capping collision risks down to the stringent 10^{-7} threshold without penetrating the covariance-defined danger volumes. Mean ΔV scaled from 14.11 mm/s (for a 10^{-4} PoC limit) to 46.29 mm/s (for a 10^{-7} PoC limit), matching or exceeding contemporary mass efficiencies. The solver demonstrated robust speeds despite complex covariance mappings, with mean computation times between 0.203 and 0.237 seconds, and autonomously identified optimal phasing at approximately 180 degrees before TCA.

Furthermore, the operational stiffness of the baseline algorithm was rigorously stress-tested under a high-fidelity dynamic environment incorporating J_2 Earth oblateness and a worst-case atmospheric drag model. The empirical results demonstrated that although minor target boundary deviations naturally occurred due to the unmodeled dynamics, they remained operationally negligible over the short-term maneuvering window. For deterministic constraints, maximum physical deviations peaked at merely 11.76 meters against the largest 3.5 km target. Probabilistic constraints maintained an even more exceptional adherence, with only 3 out of thousands of test cases slightly exceeding a 5% relative error threshold. This confirms that the unperturbed optimization strategy safely and accurately mitigates collision risks while successfully avoiding the computational burden of propagating complex environmental perturbations.

5.1. Future Work

While the proposed methodology is highly capable, the complexity of the orbital environment presents four logical avenues for future research:

The integration of high-fidelity environmental perturbations, such as J2 non-spherical Earth gravity and atmospheric drag, is required to evaluate the solver's stability over extended maneuver lead times exceeding 1.0 orbital periods. Furthermore, these perturbations could be directly integrated into the control law derivation itself, which becomes especially critical for cases in which the maneuver starts several orbits or days before TCA.

The framework must be expanded to address multi-encounter optimization, allowing satellites to navigate cascading conjunctions by defining simultaneous B-plane constraints across two or more TCAs. Concurrently, integrating dynamic covariance propagation will ensure that statistical constraints actively update as the 100% throttle thrust profile structurally alters the spacecraft's positional uncertainty.

Advancing the algorithm's initialization through machine learning remains a vital objective. Future efforts should explore Physics-Informed Neural Networks (PINNs). By embedding Hamiltonian dynamics directly into the loss function, a PINN could replace the current bisection method and provide an instantaneous warm start, potentially reducing root-finding iterations to near zero.

Embedding the algorithm onto flight-equivalent avionics for Hardware-in-the-Loop testing is essential. Validating its performance against realistic onboard memory limitations is the definitive next step toward achieving fully autonomous flight implementation on next-generation satellite platforms.

Investigating the impact of maneuver execution errors presents a critical avenue for future research. While the current probabilistic framework rigorously accounts for initial Orbit Determination (OD) uncertainties via the combined B-plane covariance matrix, the trajectory propagation assumes ideal thruster execution. Future iterations should incorporate stochastic optimal control techniques or extensive Monte Carlo simulations to evaluate the robustness of the derived bang-bang control laws against realistic onboard uncertainties, such as thrust pointing misalignments and thrust magnitude execution errors.

Bibliography

[1] 18th and 19th Space Defense Squadrons. Spaceflight Safety Handbook for Satellite Operators: 18 & 19 SDS Processes for On-Orbit Conjunction Assessment & Collision Avoidance. Technical report, Combined Force Space Component Command, 4 2023. URL `https://www.space-track.org/documents/SFS_Handbook_For_Operators_V1.7.pdf`. Version 1.7.

[2] A. Abraham. Gps transponders for space traffic management. Technical report, The Aerospace Corporation, 4 2018.

[3] A. Adriano, K. A. Scott, H. Oqab, G. Dietrich, and N. L. Azad. Neural networks for space debris classification. *Expert Systems*, 43(1):e70180, 2026.

[4] AIAA Space Traffic Management Working Group. Space traffic management (stm): Balancing safety, innovation, and growth. Institute position paper, American Institute of Aeronautics and Astronautics, 11 2017.

[5] S. Alfano and D. Oltrogge. Probability of collision: Valuation, variability, visualization, and validity. *Acta Astronautica*, 148:301–316, 2018.

[6] R. Armellin. Collision avoidance maneuver optimization with a multiple-impulse convex formulation. *Acta Astronautica*, 186:347–362, 2021.

[7] R. Bevilacqua, M. Kumar, T. Alfriend, H. Krag, and L. Anselmo. Special issue on space situational awareness from the 1st international academy of astronautics conference on space situational awareness or icssa 2017. *Acta Astronautica*, 155: 367–368, 2 2019.

[8] C. Bombardelli and J. Hernando-Ayuso. Optimal impulsive collision avoidance in low earth orbit. *Journal of Guidance, Control, and Dynamics*, 38(2):217–225, 2 2015.

[9] C. Bombardelli, J. Hernando Ayuso, and R. García Pelayo. Collision avoidance maneuver optimization. In *24th AAS/AIAA Space Flight Mechanics Meeting*, Santa Fe, New Mexico, United States, January 26–30, 2014. American Astronautical Society. Paper AAS 14-335.

[10] X. Cao, X. Ning, Z. Wang, S. Liu, F. Cheng, W. Li, and X. Lian. Intelligent sequential multi-impulse collision avoidance method for non-cooperative spacecraft based on an improved search tree algorithm. *Chinese Journal of Aeronautics*, 38(4):103205, 2025.

[11] E. Y. Choueiri. A critical history of electric propulsion: The first 50 years (1906–1956). *Journal of Propulsion and Power*, 20(2):193–203, 2004.

[12] C. Colombo, A. De Vittori, M. Omodei, J. L. Gonzalo, M. Maestrini, P. Di Lizia, P. G. Padreny, M. T. Ribell, A. G. Torrego, D. E. Antón, and R. Armellin. Sensitivity analysis of collision avoidance manoeuvre with low thrust propulsion. In *Aerospace Europe Conference 2023 – Joint 10th EUCASS & 9th CEAS*, Lausanne, Switzerland, July 9–13, 2023. EUCASS Association. EUCASS2023-093.

[13] Consultative Committee for Space Data Systems. Recommendation for space data system standards: Conjunction data message. Recommended Standard (Blue Book) CCSDS 508.0-B-1, CCSDS, Washington, DC, USA, 6 2013.

[14] L. Costi and D. Izzo. Simulation of active soft nets for capture of space debris, 11 2025.

[15] H. D. Curtis. *Orbital Mechanics for Engineering Students*. Butterworth-Heinemann, Oxford, United Kingdom, 3rd edition, 2014.

[16] A. De Vittori, M. F. Palermo, P. Di Lizia, and R. Armellin. Low-thrust collision avoidance maneuver optimization. *Journal of Guidance, Control, and Dynamics*, 45 (10):1815–1829, 10 2022.

[17] A. De Vittori, G. Dani, P. Di Lizia, and R. Armellin. Low-thrust collision avoidance design for leo missions with return to nominal orbit. In *33rd AAS/AIAA Space Flight Mechanics Meeting*, Austin, Texas, United States, January 15–19, 2023. American Astronautical Society. AAS 23-198.

[18] A. De Vittori, G. Dani, P. Di Lizia, and R. Armellin. Numerically efficient low-thrust collision avoidance maneuver design in geo regime with equinoctial orbital elements. In *2nd ESA NEO and Debris Detection Conference*, pages 1–10, Darmstadt, Germany, January 24–26, 2023. ESA Space Debris Office.

[19] A. De Vittori, M. Omodei, P. Di Lizia, R. Armellin, P. G. Padreny, M. T. Ribell, J. R. Antón, and D. E. Antón. Numerically efficient low-thrust fuel-optimal collision avoidance maneuvers with tangential firing. In *2023 AAS/AIAA Astrodynamics Specialist Conference*, Big Sky, Montana, United States, August 13–17, 2023. American Astronautical Society. AAS 23-215.

[20] A. De Vittori, G. Dani, A. Cantoni, P. Di Lizia, and R. Armellin. Combined collision avoidance and station-keeping maneuver optimization by collision probability inversion. *Journal of Spacecraft and Rockets*, 63(1):203–217, 2026.

[21] A. De Vittori, G. Dani Gabriele, A. Cantoni, P. Di Lizia, and R. Armellin. Combined collision avoidance and station-keeping maneuver optimization by collision probability inversion. *Journal of Spacecraft and Rockets*, 63(1):203–217, 1 2026.

[22] L. Dell'Elce, F. de Veld, and J. B. Pomet. Characterization of the minimum warning time for low-thrust collision avoidance maneuvers. In *2024 AAS/AIAA Astrodynamics Specialist Conference*, Broomfield, Colorado, United States, August 11–15, 2024. American Astronautical Society. hal-04871764.

[23] European Space Agency. Esa space environment report 2025. Technical report, ESA Space Safety Programme, 2025. URL `https://www.sdo.esoc.esa.int/publications/Space_Environment_Report_I9R1_20251021.pdf`. Accessed: 2026-03-13.

[24] European Space Agency Space Debris Office. ESA Space Environment Statistics. `https://sdup.esoc.esa.int/discosweb/statistics/`, 2026. Accessed: 2026-03-30.

[25] J. L. Gonzalo, C. Colombo, and P. di Lizia. Computationally efficient approaches for low-thrust collision avoidance activities. In *72nd International Astronautical Congress (IAC)*, Dubai, United Arab Emirates, October 25–29, 2021. International Astronautical Federation. IAC-21-A6.10-B6.5.5.

[26] J. L. Gonzalo, C. Colombo, and P. Di Lizia. Single-averaged models for low-thrust collision avoidance under uncertainties. In *73rd International Astronautical Congress (IAC)*, Paris, France, September 18–22, 2022. International Astronautical Federation. Paper IAC-22-A6.7.10.

[27] J. Hernando-Ayuso and C. Bombardelli. Low-thrust collision avoidance in circular orbits. *Journal of Guidance, Control, and Dynamics*, 44(5):983–995, 2021.

[28] J. I. Katz. Can a satellite dodge space debris? *Acta Astronautica*, 222:174–179, 2024.

[29] D. J. Kessler and B. G. Cour-Palais. Collision frequency of artificial satellites: The creation of a debris belt. *Journal of Geophysical Research*, 83(A6):2637–2646, 1978.

[30] K. Lee, C. Park, and S. Y. Park. Near-optimal guidance and control for spacecraft collision avoidance maneuvers. In *AIAA SPACE 2014 Conference and Exposition*, San Diego, California, United States, August 4–7, 2014. American Institute of Aeronautics and Astronautics. AIAA 2014-4114.

[31] D. Lev, R. M. Myers, K. M. Lemmer, J. Kolbeck, H. Koizumi, and K. Polzin. The technological and commercial expansion of electric propulsion. *Acta Astronautica*, 159:213–227, 6 2019.

[32] H. G. Lewis and G. Skelton. Safety considerations for large constellations of satellites. *Journal of Space Safety Engineering*, 11(3):439–445, 2024.

[33] T.-C. Lin and Y.-H. Lee. Review of electric propulsion in cubesat. In *2022 International Conference on Astronautics and Space Exploration (iCASE)*, Hsinchu, Taiwan, November 3–5, 2022. IEEE.

[34] J. Liu, H. Jiang, X. Yang, F. Li, and N. Zhao. Space debris research progress of china. *Chinese Journal of Space Science*, 40(5):956–961, 9 2020.

[35] W. Lu, Z. Geng, Y. Huang, Z. Mu, S. Wu, V. Y. Razoumny, and Y. N. Razoumny. Optimizing spacecraft collision avoidance maneuvers under risk uncertainty from multiple space debris. *Chinese Journal of Aeronautics*, 37:103814, 9 2025.

[36] M. Maestrini, A. De Vittori, J. L. Gonzalo Gómez, C. Colombo, P. Di Lizia, J. Míguez Arenas, M. Sanjurjo Rivo, A. Diez Martín, P. Gago Padreny, and D. Escobar Antón. Electrocam: Assessing the effect of low-thrust uncertainties on orbit propagation. In *Proceedings of the 2nd ESA NEO and Debris Detection Conference*, Darmstadt, Germany, January 24–26, 2023. ESA Space Debris Office.

[37] L. Niccolai and G. Mengali. Performance estimate of a spin-stabilized drag sail for spacecraft deorbiting. *Applied Sciences*, 14(2):612, 2024.

[38] A. Pak, J. Correa, and M. Adams. Robust joint target detection and tracking for space situational awareness. *Journal of Guidance, Control, and Dynamics*, 41(1): 119–136, 1 2018.

[39] M. F. Palermo. Numerically efficient methods for impulsive and low-thrust collision avoidance manoeuvre design. Master's thesis, Politecnico di Milano, Milan, Italy, 2021. URL `https://www.politesi.polimi.it/retrieve/28b04b10-e94f-4b79-9b3f-6ef3227117fb/MasterThesis_MFP.pdf`.

[40] S. Patnala and A. Abdin. Spacecraft collision avoidance: data management, risk assessment, decision planning models and algorithms. In *Space Data Management*, pages 15–45. Springer, Cham, Switzerland, 2024.

[41] Z. Pavanello, L. Pirovano, R. Armellin, A. De Vittori, and P. Di Lizia. A convex optimization method for multiple encounters collision avoidance maneuvers. In *AIAA SCITECH 2024 Forum*, Orlando, Florida, United States, January 8–12, 2024. Amer-

ican Institute of Aeronautics and Astronautics. doi: 10.2514/6.2024-0845. AIAA 2024-0845.

[42] Z. Pavanello, L. Pirovano, R. Armellin, A. De Vittori, and P. Di Lizia. A convex formulation for collision avoidance maneuver strategies during low-thrust phases. In *AIAA SCITECH 2024 Forum*, Orlando, Florida, United States, January 8–12, 2024. American Institute of Aeronautics and Astronautics. doi: 10.2514/6.2024-0844. AIAA 2024-0844.

[43] Z. Pavanello, L. De Maria, A. De Vittori, M. Maestrini, P. Di Lizia, and R. Armellin. Cammary: A review of spacecraft collision avoidance manoeuvre design methods. *Acta Astronautica*, 236:770–789, 2025.

[44] Z. Pavanello, L. Pirovano, and R. Armellin. Recursive polynomial method for fast collision avoidance maneuver design. *IEEE Transactions on Aerospace and Electronic Systems*, 61(2):2914–2925, 4 2025.

[45] Z. Pavanello, L. Pirovano, R. Armellin, A. De Vittori, and P. Di Lizia. Collision avoidance maneuver optimization during low-thrust propelled trajectories. *Astrodynamics*, 9(2):247–271, 2025.

[46] E. M. Polli, J. L. Gonzalo, and C. Colombo. Semi-analytical model for autonomous fuel-optimal low-thrust collision avoidance with eclipse constraints. In *35th AAS/AIAA Space Flight Mechanics Meeting*, Kauai, Hawaii, United States, January 19–23, 2025. American Astronautical Society. TODO: Add AAS paper number.

[47] S. S. Popova. « » [legal gaps in the definition of space debris]. *Moscow Journal of International Law*, (2):39–48, 2025.

[48] J. A. Pulido Cobo, N. Sánchez Ortiz, I. Grande Olalla, and K. Merz. Coram: Esa's collision risk assessment and avoidance manoeuvres computation tool. In *2nd International Academy of Astronautics Conference on Dynamics and Control of Space Systems (DyCoSS)*, Rome, Italy, March 24–26, 2014. International Academy of Astronautics. Paper IAA-AAS-DyCoSS2-05-03.

[49] T. Pultarova. Spacex starlink satellites caused 25,000 near-misses in just 6 months. Space.com, 7 2023. URL `https://www.space.com/starlink-satellite-conjunction-increase-threatens-space-sustainability`. Accessed: 2026-04-30.

[50] M. Rasotto, R. Armellin, and P. Di Lizia. Multi-step optimization strategy for fuel-optimal orbital transfer of low-thrust spacecraft. *Engineering Optimization*, 48(3): 519–542, 2016.

[51] J. A. Reiter and D. B. Spencer. Solutions to rapid collision-avoidance maneuvers constrained by mission performance requirements. *Journal of Spacecraft and Rockets*, 55(4):1039–1047, 2018.

[52] A. Schiavo. Numerically efficient design of low-thrust collision avoidance manoeuvres. Master's thesis, Politecnico di Milano, 2020.

[53] M. Sorge, T. Harris, C. Colombo, M. Hejduk, N. Fitz-Coy, R. Sheppard, E. Kerr, U. Dasgupta, N. Berend, D. Escobar, M. Michel, and C. Perez. Space traffic management: Improvements to spacecraft collision avoidance (cola). *Acta Astronautica*, 229:600–605, 4 2025.

[54] A. Stoev, P. Maglova, O. Ognyanov, and A. Radev. "space debris" and its impact on the space activity of civilization and the ecology of earth. In *Near Space – A Common Goal Conference*, Veliko Tarnovo, Bulgaria, November 12–14, 2025. Vasil Levski National Military University. doi: 10.34660/INF.2025.51.30.045.

[55] R. Ströbel and E. Stoll. Analytical assessment of short-duration low-thrust collision-avoidance maneuvers. *Journal of Guidance, Control, and Dynamics*, 47(2):358–365, 2 2024.

[56] E. Taheri and J. L. Junkins. Generic smoothing for optimal bang-off-bang spacecraft maneuvers. *Journal of Guidance, Control, and Dynamics*, 41(11):2467–2472, 2018.

[57] G. Tang, F. Jiang, and J. Li. Fuel-optimal low-thrust trajectory optimization using indirect method and successive convex programming. *IEEE Transactions on Aerospace and Electronic Systems*, 54(4):2053–2066, 2018.

[58] Y. Wang and F. Topputo. Indirect optimization of fuel-optimal many-revolution low-thrust transfers with eclipses. *IEEE Transactions on Aerospace and Electronic Systems*, 59(1):39–51, 2 2023.

[59] Z. Wang and M. J. Grant. Minimum-fuel low-thrust transfers for spacecraft: A convex approach. *IEEE Transactions on Aerospace and Electronic Systems*, 54(5): 2274–2290, 2018.

[60] J. R. Wertz, editor. *Spacecraft Attitude Determination and Control*, volume 73 of *Astrophysics and Space Science Library*. Springer, Dordrecht, Netherlands, 1978.

[61] R. Wolff, D. Poster, M. Borowitz, B. Walling, A. Parikka, E. Baumann, C. Diamantopoulos, Q. Collins, J. Lorentz, S. Pritchard, et al. Conjunction data message (cdm) specification for traffic coordination system for space (tracss). Technical report, National Oceanic & Atmospheric Administration (NOAA) Office of Space Commerce, 4

2025. URL https://space.commerce.gov/wp-content/uploads/2025/04/Conjunction-Data-Message-CDM-Specification-for-Traffic-Coordination-System-for-Space-Draft-TraCSS-Spec-001.pdf. Draft TraCSS-Spec-001 Release 2.0.

[62] D. Wright. Colliding satellites: Consequences and implications. Technical report, Union of Concerned Scientists, 2 2009. URL https://www.ucsusa.org/sites/default/files/2019-10/SatelliteCollision-2-12-09.pdf.

[63] J. Zhou and K. Li. Multi-object collision avoidance method based on sequential convex optimization. *Chinese Journal of Space Science*, 46(1):189–197, 2026.

List of Figures

1.1 Overview of space debris distribution around the Earth [23] 1
1.2 Evolution of the launch traffic near LEO per mission type in object number [23] 2
1.3 Number of cumulative collisions in LEO in the simulated scenarios of long-term evolution of the environment. 2
1.4 Historical trend of number of fragments produced by fragmentation events. The last bin covers until the end of 2024 [23] 3
1.5 Number of conjunction risk mitigation maneuvers performed by satellites from the Starlink constellation in six-monthly periods covering December 2020 through May 2024. The linear model has a gradient of approximately 40 maneuvers per day [32] 3
1.6 Dodging a passive object to keep the satellite safe (ESA Video) 4
1.7 Diagram of the thesis content 9

2.1 Visualization of ECI and RTN frames[37] 14
2.2 Encounter frame and B-plane: snapshot of $O_p - O_s$ encounter geometry[16] 16
2.3 Visualization of combined cross-sectional radius s_A[39], where the centers of the two objects are depicted as red dots. 20
2.4 Visualization of covariances intersecting along orbital paths [2] 21
2.5 Conjunction encounter and covariance 23
2.6 3D Combined Covariance Ellipsoid 24
2.7 Geometric and probabilistic constraints 25
2.8 Normalized B-plane heatmaps 26

3.1 Diagram showing the rationale behind the method 31
3.2 Simplified block diagram of the Differential Correction method. 37
3.3 Velocity costate profile comparison. The energy-optimal CAM (dashed lines) closely mirrors the true fuel-optimal control (solid lines) initiated two orbits before TCA [16] 38

3.4 Qualitative visualization of the STM propagating an instantaneous perturbation induced by the propulsion unit forward to the terminal state. 40
3.5 Effect of incrementing the ignition time t_{on}, visually demonstrating the removal of a slice of the active thrust profile. 41
3.6 Effect of incrementing the cutoff time t_{off}, visually demonstrating the addition of a slice to the active thrust profile. 41
3.7 Complete workflow of the proposed methodology, highlighting the transition from the analytical energy-optimal initialization (Phase 1) to the numerical differential correction loop (Phase 2). 44
3.8 Block diagram of the Bisection Method for Velocity Costate Threshold . . 45

4.1 Violin plots for 0.5 km 55
4.2 Violin plots for 1.5 km 56
4.3 Violin plots for 2.5 km 56
4.4 Violin plots for 3.5 km 56
4.5 Preferred firing times 57
4.6 CDF of miss distance constraint violations 58
4.7 Visualization of the B-plane comparing the satellite positions before and after the maneuver 59
4.8 Costate vector evolution 60
4.9 Acceleration profile 60
4.10 Violin plots for PoC 10^{-4} 63
4.11 Violin plots for PoC 10^{-5} 63
4.12 Violin plots for PoC 10^{-6} 63
4.13 Violin plots for PoC 10^{-7} 64
4.14 Preferred firing times for PoC targets 65
4.15 CDF of PoC constraint violations 66
4.16 Evolution of the costate vector components over time for sample ID 1 under PoC constraints. 67
4.17 Acceleration components over time for sample ID 1 under PoC constraints. 67
4.18 Visualization of the B-plane comparing satellite positions before and after the maneuver, showcasing the dual nature of the constraint: targeting a specific PoC physically manifests as an elliptical Squared Mahalanobis Distance boundary 67
4.19 Perturbation effects validation diagram 68
4.20 Comparison of perturbed and unperturbed miss distances (0.5 km and 1.5 km) 70

4.21 Comparison of perturbed and unperturbed miss distances (2.5 km and 3.5 km) 71
4.22 Comparison of perturbed and unperturbed PoC values (10^{-4} and 10^{-5}) . . 72
4.23 Comparison of perturbed and unperturbed PoC values (10^{-6} and 10^{-7}) . . 72

List of Tables

4.1 Mean ΔV for MD thresholds 55
4.2 Mean ΔV for PoC thresholds 64

List of Symbols

Variable	Description	SI unit
$\boldsymbol{r}$	Spacecraft position vector	m
$\boldsymbol{v}$	Spacecraft velocity vector	m/s
m	Spacecraft mass	kg
μ	Earth's standard gravitational parameter	m^3/s^2
T_{max}	Maximum available thrust	N
I_{sp}	Engine's specific impulse	s
g_0	Standard gravitational acceleration	m/s^2
$\boldsymbol{a}_c$	Commanded control acceleration vector	m/s^2
a_{max}	Constant maximum acceleration	m/s^2
$\boldsymbol{\alpha}$	Thrust pointing unit vector	-
ϵ or u	Dimensionless throttle factor	-
Δv	Total required velocity change	m/s
$\boldsymbol{b}_f$	B-plane relative position vector	m
s_A	Combined cross-sectional radius	m
$\boldsymbol{C}$	Combined B-plane covariance matrix	m^2
PoC	Probability of Collision	-
d_M^2	Squared Mahalanobis distance	-
t_0	Initial maneuver epoch	s
t_f	Time of Closest Approach (TCA)	s
t_{on}	Engine ignition time	s
t_{off}	Engine cutoff time	s
$\boldsymbol{\lambda}_r$	Position costate vector	s/m
$\boldsymbol{\lambda}_v$	Velocity costate vector	s^2/m
$\boldsymbol{\Phi}(t, t_0)$	State Transition Matrix (STM)	-

Acknowledgements

This work is the result of more than 1 year of dedication, successes and failures. The outcome would not have been possible without the precious help and guidance of my two supervisors: Prof. Pierluigi Di Lizia and Prof. Haijun Rong. Each of them provided me with key tips and knowledge, stimulating my critical thinking and giving me room to explore solutions on my own.

I extend my thanks to the committee members for their time dedicated to reading and evaluating my work.

Special thanks go to my family. Thanks Mom and Dad for your support during this beautiful but challenging chapter of my life, for standing by every choice that brought me here. You raised me as a free soul, never imposing your worldviews, instead teaching me to find answers through self-discipline and freedom. Thanks to my siblings Luca, Alessandro, Elisa and Matteo for your presence in my heart and mind during all these years, most of which I had to spend apart from you all. I wish we could be closer and I cannot wait to see you again.

Many thanks to Gianluca for being such a great companion during the Covid-19 pandemic. We really motivated each other and reached great goals together.

Thanks are owed to Danylo for sharing his living room as a bedroom with me during the hardest time of my student career, beside being my long time irreplaceable friend.

I owe a deep debt of gratitude to Francesco and Raffaele for being such good colleagues, flatmates, and friends during this life-changing adventure in China.

And finally, a broad thanks to all the people I have met along the way, for your contribution to shaping my personality and broadening my view of the World. From Italy, through Mexico, ending up to China: **grazie mille, muchísimas gracias, 感谢**!