

UNIFIED COVECTOR MAPPING FRAMEWORK FOR LOW-THRUST TRAJECTORY OPTIMIZATION

Niccolò Michelotti ^{*} and Francesco Topputo [†]

Low-thrust trajectory optimization often requires balancing the robustness and computational efficiency of direct and indirect methods while accounting for the influence of different state coordinate representations. This paper introduces a Unified Covector Mapping Framework that establishes analytical relationships between continuous-time indirect costates with discrete nonlinear programming Lagrange multipliers across different coordinate charts. These relationships enable a hybrid autonomous guidance architecture that transfers optimal solutions between direct and indirect solvers and across multiple coordinate representations (e.g., Modified Equinoctial Elements and Cartesian coordinates). This framework mitigates the disadvantages of individual approaches and enables efficient hybrid optimization strategies to solve challenging optimal control problems. Following verification of the proposed costate mapping, the hybrid algorithm performance is evaluated using 450 interplanetary transfer scenarios, demonstrating a 9% increase in convergence rate, a 4.4% reduction in propellant consumption, and a 46% decrease in computational time compared with a standard indirect shooting method. These results demonstrate that the proposed framework improves both robustness and computational efficiency, supporting its application to onboard autonomous trajectory optimization.

1 INTRODUCTION

Low-thrust electric propulsion has become a critical enabling technology for the next generation of deep-space missions due to its high specific impulse and mass efficiency. However, the prolonged thrusting arcs dictate that trajectory optimization must be treated as a continuous optimal control problem.^{1,2} Traditionally, these trajectories are generated on the ground by teams of mission analysts. These human-in-the-loop operations are expensive and require large ground infrastructures to maintain frequent communication windows with the probe. To sustain the momentum of deep-space exploration, especially with miniaturized platforms, a paradigm shift toward self-driving spacecraft is necessary. To address this challenge, the ERC-funded EXTREMA (Engineering Extremely Rare Events in Astrodynamics for Deep-Space Missions in Autonomy) project aims to enable self-driving interplanetary CubeSats capable of traveling in deep space without requiring interaction from the ground.³ Within this framework, autonomous guidance plays a central role, as the spacecraft must be able to recompute feasible and efficient low-thrust trajectories onboard, using the current navigation estimate and without support from ground operators.⁴

Spacecraft trajectory optimization relies heavily on the accurate numerical solution of nonlinear optimal control problems (OCPs). The state of the art generally classifies these solution techniques

^{*}PhD Student, Department of Aerospace Science and Technology, Politecnico di Milano, via La Masa, 34, Milan, 20156, Italy. niccolo.michelotti@polimi.it

[†]Full Professor, Department of Aerospace Science and Technology, Politecnico di Milano, via La Masa, 34, Milan, 20156, Italy. francesco.topputo@polimi.it. AIAA Senior Member.

into direct and indirect methods.^{5,6} Direct methods transcribe the continuous-time OCP into a finite-dimensional nonlinear programming (NLP) problem.⁷ They offer relatively wide convergence basins and handle path constraints well, but suffer from high computational burdens and struggle to accurately capture discontinuous bang-bang control structures without severe mesh refinement. Conversely, indirect methods apply Pontryagin’s Maximum Principle (PMP) to formulate a two-point boundary value problem (TPBVP). While computationally fast and theoretically precise, they are characterized by extremely narrow convergence radii, requiring highly accurate initial guesses for the non-intuitive dual variables (i.e., costates). Concurrently, a parallel challenge in trajectory optimization is the appropriate choice of coordinate representations. The dynamical regularity of the chosen reference frame heavily influences the convergence performance of the associated solver.^{8,9} For example, formulating the equations of motion in Cartesian coordinates yields simpler algebraic forms for path constraints (such as solar eclipses) and environmental perturbations, making them favorable for direct NLP transcriptions. Conversely, the use of Modified Equinoctial Elements (MEE) regularizes Keplerian dynamics, significantly enhancing the convergence for multi-revolution transfers, albeit at the cost of highly complex state-costate differential equations.

Previous literature has explored the possibility of linking different low-thrust optimization problem formulations and their solutions.^{10,11} The Covector Mapping Theorem (CMT)¹² provides a theoretical link between discrete NLP multipliers and continuous indirect costates within the same coordinate frame. Separately, variational approaches¹³ have been derived to map costates across different coordinate charts within continuous indirect formulations. However, practical autonomous guidance architectures require the flexibility to do both simultaneously: an onboard computer may need to initialize a problem in one formulation (e.g., indirect MEE) to leverage dynamic regularity, and refine the solution in an entirely different formulation (e.g., direct Cartesian) to efficiently handle complex operational constraints. Therefore, a critical gap in the existing literature is the absence of a unified theoretical framework that bridges both the solver transcription and the coordinate geometry. This paper addresses this gap by developing a Unified Covector Mapping Framework for nonlinear low-thrust trajectory optimization. By formally composing the discrete-to-continuous Covector Mapping Theorem with the variational coordinate adjoint transformation, we establish a mathematically consistent bridge to map dual variables across both transcription methods and coordinate charts. This theoretical foundation is subsequently translated into robust, flexible, and computationally efficient hybrid guidance algorithms tailored for onboard execution. By treating dual information as a transportable mathematical asset, the proposed pipeline leverages direct-indirect warm starts, objective function continuation, and Cartesian-MEE dual transfers to systematically mitigate the limitations of individual solvers.

The remainder of this paper is organized as follows. Section 2 formulates the low-thrust optimal control problem and defines the spacecraft dynamics. Section 3 derives the theoretical Unified Covector Mapping Framework and outlines the proposed hybrid guidance algorithms. Section 4 presents numerical results validating the mapping between approaches and reference frames, followed by a large-scale performance analysis carried out across 450 interplanetary scenarios to Venus and Mars. Finally, Section 5 concludes the paper.

2 PROBLEM FORMULATION

2.1 The Continuous Optimal Control Problem

A general optimal control problem for low-thrust trajectory optimization requires determining the state history $\mathbf{y}(t) \in \mathbb{R}^n$ and the control history $\mathbf{u}(t) \in \mathbb{R}^m$ that minimize a prescribed performance

index. This is mathematically formulated as minimizing the Bolza cost functional:

$$J = \Phi(\mathbf{y}(t_0), t_0, \mathbf{y}(t_f), t_f) + \int_{t_0}^{t_f} L(\mathbf{y}(t), \mathbf{u}(t), t) dt, \quad (1)$$

subject to the continuous nonlinear dynamics constraints,

$$\dot{\mathbf{y}}(t) = \mathbf{f}(\mathbf{y}(t), \mathbf{u}(t), t), \quad (2)$$

the initial and terminal boundary conditions,

$$\phi(\mathbf{y}(t_0), t_0, \mathbf{y}(t_f), t_f) = \mathbf{0}, \quad (3)$$

and the inequality path constraints (e.g., bounds on control or state-dependent operational limits),

$$\mathbf{C}(\mathbf{y}(t), \mathbf{u}(t), t) \leq \mathbf{0}. \quad (4)$$

Depending on the specific mission objectives, the Lagrangian integrand L and the terminal cost Φ are adapted to represent different performance criteria. To minimize propellant consumption, a fuel-optimal objective is defined as¹⁴

$$J = \frac{1}{T_{max} I_{sp} g_0} \int_{t_0}^{t_f} \mathbf{u}(t) dt, \quad (5)$$

where I_{sp} is the specific impulse, g_0 is the Earth gravity acceleration, $\mathbf{u}$ is the control logic, and $[t_0, t_f]$ is the time interval of the simulation. The minimization of propellant consumption yields a discontinuous bang-bang control structure that can be cumbersome for solver convergence. Conversely, setting the objective function as¹⁴

$$J = \frac{1}{T_{max} I_{sp} g_0} \int_{t_0}^{t_f} \mathbf{u}(t)^2 dt \quad (6)$$

defines an energy-optimal problem, which yields a smooth, continuous thrust profile that is significantly easier to solve computationally. The transition between these two formulations is a core mechanism of the hybrid algorithms proposed in this work. Lastly, for time-optimal problems no integral term is present, and the objective function reduces to minimizing the final time

$$J = t_f. \quad (7)$$

To ensure numerical stability across all formulations and solvers employed, the equations of motion are strictly non-dimensionalized using canonical units for length, time, and mass.

2.2 Spacecraft Dynamics and Coordinate Representations

The choice of the state vector $\mathbf{y}$ dictates the algebraic structure of the vector field $\mathbf{f}$. This work leverages two primary coordinate representations: Cartesian coordinates and Modified Equinoctial Elements (MEE).

Let $\mathbf{x} = [\mathbf{r}^\top, \mathbf{v}^\top, m]^\top \in \mathbb{R}^7$ denote the state vector in the Cartesian frame, where $\mathbf{r}$ is position, $\mathbf{v}$ is velocity, and m is the spacecraft mass. The equations of motion are expressed as:

$$\dot{\mathbf{x}} = \begin{bmatrix} \mathbf{v} \\ -\frac{\mu}{r^3} \mathbf{r} + \frac{T_{max}}{m} u \hat{\boldsymbol{\alpha}}_I \\ -\frac{T_{max}}{I_{sp} g_0} u \end{bmatrix}, \quad (8)$$

where μ is the gravitational parameter, T_{max} is the maximum available thrust, I_{sp} is the specific impulse, g_0 is standard gravity, $u \in [0, 1]$ is the thrust throttle, and $\hat{\alpha}_I$ is the thrust direction unit vector parameterized in the Inertial reference frame.

Alternatively, let $\mathbf{z} = [p, f, g, h, k, L, m]^\top \in \mathbb{R}^7$ denote the state vector augmented with mass in the MEE frame. The MEE variables resolve the singularities associated with classical orbital elements at zero eccentricity and zero inclination. The governing dynamics in MEE take the form of Gauss's variational equations:

$$\dot{\mathbf{z}} = \begin{bmatrix} \mathbf{A}(\mathbf{z}) + \mathbf{B}(\mathbf{z}) \frac{T_{max}}{m} u \hat{\alpha}_{RTN} \\ -\frac{T_{max}}{I_{sp} g_0} u \end{bmatrix}, \quad (9)$$

where the unperturbed Keplerian drift vector $\mathbf{A}(\mathbf{z})$ and the control multiplier matrix $\mathbf{B}(\mathbf{z})$ are defined as:

$$\mathbf{A}(\mathbf{z}) = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \sqrt{\frac{\mu}{p^3}} w^2 \end{bmatrix}, \quad \mathbf{B}(\mathbf{z}) = \sqrt{\frac{p}{\mu}} \begin{bmatrix} 0 & \frac{2p}{w} & 0 \\ \sin L & \frac{1}{w} [(w+1) \cos L + f] & -\frac{g}{w} (h \sin L - k \cos L) \\ -\cos L & \frac{1}{w} [(w+1) \sin L + g] & \frac{f}{w} (h \sin L - k \cos L) \\ 0 & 0 & \frac{s^2 \cos L}{2w} \\ 0 & 0 & \frac{s^2 \sin L}{2w} \\ 0 & 0 & \frac{1}{w} (h \sin L - k \cos L) \end{bmatrix} \quad (10)$$

with the auxiliary variables $w = 1 + f \cos L + g \sin L$ and $s^2 = 1 + h^2 + k^2$. Notably, Gauss's variational equations dictate that the thrust direction unit vector $\hat{\alpha}_{RTN}$ must be parameterized in the moving Radial-Transverse-Normal (RTN) frame. Transitioning between these formulations requires a continuous, invertible mapping ρ , such that $\mathbf{z}(t) = \rho(\mathbf{x}(t))$. Let

$$\alpha = h^2 - k^2, \quad w = 1 + f \cos L + g \sin L, \quad s^2 = 1 + h^2 + k^2, \quad M = \frac{1}{s^2 \sqrt{p}}.$$

The nonlinear transformation between the MEE and the Cartesian state is defined as¹³

$$\mathbf{x}(t) = \rho(\mathbf{z}(t)) = \begin{bmatrix} \frac{p}{ws^2} (\cos L + \alpha \cos L + 2hk \sin L) \\ \frac{p}{ws^2} (\sin L - \alpha \sin L + 2hk \cos L) \\ \frac{2p}{ws^2} (h \sin L - k \cos L) \\ -M [\sin L + \alpha \sin L - 2hk \cos L + g - 2fhk + \alpha g] \\ -M [-\cos L + \alpha \cos L + 2hk \sin L - f + 2ghk + \alpha f] \\ 2M (h \cos L + k \sin L + fh + gk) \end{bmatrix}. \quad (11)$$

The fundamental geometric Jacobian of this transformation, necessary for mapping dual variables, is defined as $\mathbf{J}_\rho(\mathbf{z}) = \frac{\partial \rho(\mathbf{z})}{\partial \mathbf{z}}$.

For the specific case of an interplanetary rendezvous mission, the generic boundary conditions in Eq. (3) take a specific algebraic form depending on the chosen coordinate representation. In Cartesian coordinates, a rendezvous requires matching the target's position $\mathbf{r}_{tgt}$ and velocity $\mathbf{v}_{tgt}$

at the final time t_f , while the initial state is fixed to the departure conditions. Thus, the boundary constraint vector is defined as:

$$\phi_{Car} = \begin{bmatrix} \mathbf{r}(t_0) - \mathbf{r}_0 \\ \mathbf{v}(t_0) - \mathbf{v}_0 \\ m(t_0) - m_0 \\ \mathbf{r}(t_f) - \mathbf{r}_{tgt}(t_f) \\ \mathbf{v}(t_f) - \mathbf{v}_{tgt}(t_f) \end{bmatrix} = \mathbf{0}. \quad (12)$$

Alternatively, in MEE coordinates, the rendezvous condition requires the spacecraft's orbital elements to match the target's elements. Letting $\alpha = [p, f, g, h, k, L]^\top$ denote the orbital element subset of $\mathbf{z}$, the boundary constraints are formulated as:

$$\phi_{MEE} = \begin{bmatrix} \alpha(t_0) - \alpha_0 \\ m(t_0) - m_0 \\ \alpha(t_f) - \alpha_{tgt}(t_f) \end{bmatrix} = \mathbf{0}. \quad (13)$$

While the physical objective of the mission remains identical, the mathematical formulation of these transversality conditions significantly influences the complexity of deriving the terminal costates in indirect methods, making the ability to dynamically map between formulations highly advantageous.

2.3 Legendre-Gauss-Radau Pseudospectral Transcription

To numerically solve the continuous OCP via direct methods, the problem must be transcribed into a finite-dimensional NLP. This work employs the Legendre-Gauss-Radau (LGR) pseudospectral method, chosen for its quasi-exponential convergence rate on smooth intervals and its mathematically rigorous adherence to the Covector Mapping Theorem.¹⁵ In addition, local LGR collocation schemas significantly increase the sparsity of the NLP Jacobian matrices, thereby reducing the overall computational effort.

The LGR discretization defines a set of collocation points $(\tau_1, \dots, \tau_N)$ on the normalized time domain $[-1, +1]$. These points correspond to the roots of $P_{N-1}(\tau) + P_N(\tau) = 0$, where P_N is the Legendre polynomial of degree N . An additional non-located terminal point $\tau_{N+1} = +1$ is included to capture the final boundary conditions. The physical time domain $t \in [t_0, t_f]$ is mapped to the normalized domain via the affine transformation:

$$t = \frac{t_f - t_0}{2}\tau + \frac{t_f + t_0}{2}. \quad (14)$$

The continuous state and control variables are approximated by summations of Lagrange polynomials $L_i(\tau)$. Because the optimal control is not explicitly required at the final boundary τ_{N+1} in the LGR scheme, the state approximation utilizes $N + 1$ points while the control approximation utilizes N points:

$$\mathbf{y}(\tau) \approx \mathbf{Y}(\tau) = \sum_{i=1}^{N+1} \mathbf{Y}_i L_i(\tau), \quad \mathbf{u}(\tau) \approx \mathbf{U}(\tau) = \sum_{i=1}^N \mathbf{U}_i L_i(\tau). \quad (15)$$

By differentiating the state polynomial approximation, the continuous dynamics (Eq. 2) are replaced by a set of discrete algebraic equality constraints evaluated strictly at the N collocation

nodes:

$$\sum_{i=1}^{N+1} D_{ki} \mathbf{Y}_i - \frac{t_f - t_0}{2} \mathbf{f}(\mathbf{Y}_k, \mathbf{U}_k, \tau_k) = \mathbf{0}, \quad k = 1, \dots, N, \quad (16)$$

where $D_{ki} = \dot{L}_i(\tau_k)$ are the elements of the $N \times (N + 1)$ Radau pseudospectral differentiation matrix $\mathbf{D}$.

Finally, the continuous integral term of the Bolza cost functional (Eq. 1) is evaluated using the LGR Gaussian quadrature rule:

$$J \approx \Phi(\mathbf{Y}_1, \tau_1, \mathbf{Y}_{N+1}, \tau_{N+1}) + \frac{t_f - t_0}{2} \sum_{k=1}^N w_k L(\mathbf{Y}_k, \mathbf{U}_k, \tau_k), \quad (17)$$

where w_k are the highly specific quadrature weights associated with each collocation node. This rigorous transcription converts the continuous time OCP into a sparse NLP, paving the way for the extraction of the dual KKT multipliers in the unified mapping framework.

3 COVECTOR MAPPING FRAMEWORK

This section establishes the theoretical foundation of the Unified Covector Mapping Framework. First, the traditional Covector Mapping Theorem (CMT) is reviewed, detailing the bridge between discrete and continuous dual variables within a single coordinate chart. Next, the variational approach for mapping continuous costates across different geometric representations is derived. Finally, these two independent operators are mathematically composed to formalize the unified mapping framework.

3.1 Theory of the Covector Mapping Principle

The original continuous time OCP is denoted as Problem P and it can be solved through two distinct theoretical pathways: indirect and direct methods. By deriving the continuous optimality conditions and applying Pontryagin's Minimum Principle, Problem P is transformed into a two-point boundary value problem (Problem P^λ). This dualization process introduces the costates (Lagrange multipliers associated with the continuous dynamics). Solving P^λ numerically via discretization (Problem $P^{\lambda N}$) constitutes the indirect approach.

Conversely, the direct approach first discretizes Problem P over a temporal grid to form a finite-dimensional Nonlinear Programming problem (Problem P^N). The NLP is then solved using numerical optimization solvers, yielding a set of discrete Karush-Kuhn-Tucker (KKT) multipliers (Problem $P^{N\lambda}$). In general, the operations of discretization and dualization are not commutative, meaning the discrete KKT multipliers of $P^{N\lambda}$ do not natively equal the continuous costates of $P^{\lambda N}$.

The Covector Mapping Theorem (CMT) mathematically closes this non-commutation gap.¹² A schematic of this principle, illustrating the commutative relationship between the direct and indirect approaches, is shown in Figure 1.

Pseudospectral (PS) collocation schemes, particularly those utilizing Legendre-Gauss-Radau (LGR) nodes, have been proven to satisfy the CMT. By analyzing the KKT stationarity conditions of the LGR-discretized NLP, the precise relationships between the discrete multipliers and the continuous physical costates can be extracted.¹⁵ For an LGR PS method defined over N collocation points, the




Figure 1. Schematic of the Covector Mapping Principle.¹² Red arrows denote operations in the indirect approach, while blue arrows denote the direct approach.

exact mapping is:

$$\tilde{\lambda}_{N+1} = \mathbf{D}_{N+1}^\top \Lambda \quad (18)$$

$$\tilde{\lambda}_k = \frac{\Lambda_k}{w_k}, \quad 1 \leq k \leq N \quad (19)$$

$$\tilde{\gamma}_k = \frac{\Gamma_k}{w_k}, \quad 1 \leq k \leq N \quad (20)$$

$$\tilde{\psi} = \Psi \quad (21)$$

The left-hand sides represent the mapped continuous-time Lagrange multipliers: $\tilde{\lambda}_k$ for the dynamics (costates), $\tilde{\gamma}_k$ for path constraints, and $\tilde{\psi}$ for boundary constraints. The right-hand sides denote the raw KKT multipliers output by the NLP solver (Λ_k , Γ_k , Ψ). Here, w_k and $\mathbf{D}_{N+1}$ are the LGR quadrature weights and the terminal column of the differentiation matrix, respectively, as derived in Section 2.3.

3.2 Variational Coordinate Costate Mapping

A variational approach can be used to map the costates among different coordinate frames.¹³ Let $\mathbf{x} \in \mathbb{R}^n$ denote the primary state vector (e.g., Cartesian coordinates) with its associated control vector $\mathbf{u}$, and let $\mathbf{z} \in \mathbb{R}^n$ denote the alternative state vector (e.g., MEEs) with its corresponding control vector α .

The nonlinear, one-to-one, continuous, and invertible mapping between the two coordinate charts is defined by ρ :

$$\mathbf{z}(t) = \rho(\mathbf{x}(t)). \quad (22)$$

The dynamics in the primary chart $\mathbf{x}$ are given by $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t)$, while the transformed differential equations in the alternative chart $\mathbf{z}$ are denoted by $\dot{\mathbf{z}} = \mathbf{g}(\mathbf{z}(t), \mathbf{w}(t), t)$.

Assuming that time does not appear explicitly in the mapping transformation, taking the time derivative of Eq. (22) yields the kinematic relationship between the vector fields:

$$\dot{\mathbf{z}}(t) = \left[\frac{\partial \rho(\mathbf{x}(t))}{\partial \mathbf{x}} \right] \dot{\mathbf{x}}(t) \implies \mathbf{g} = \mathbf{J}_\rho(\mathbf{x}) \mathbf{f}, \quad (23)$$

where $\mathbf{J}_\rho(\mathbf{x}) = \frac{\partial \rho(\mathbf{x})}{\partial \mathbf{x}} \in \mathbb{R}^{n \times n}$ is the fundamental Jacobian matrix of the coordinate mapping evaluated at $\mathbf{x}(t)$.

Let $\boldsymbol{\lambda}_x$ and $\boldsymbol{\lambda}_z$ denote the costate vectors associated with the states $\mathbf{x}$ and $\mathbf{z}$, respectively. The optimal control Hamiltonians for the two representations are defined as:

$$H_x = L(\mathbf{x}(t), \mathbf{u}(t), t) + \boldsymbol{\lambda}_x^\top \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t), \quad (24)$$

$$H_z = L(\mathbf{z}(t), \mathbf{w}(t), t) + \boldsymbol{\lambda}_z^\top \mathbf{g}(\mathbf{z}(t), \mathbf{w}(t), t). \quad (25)$$

Because the Euler-Lagrange equations are strictly satisfied along extremal trajectories regardless of the chosen coordinate representation, the Hamiltonian remains invariant under any coordinate transformation for motion referenced to the same frame. Setting $H_x = H_z$ and noting that the physical Lagrangian L remains identical ($L(\mathbf{x}, \mathbf{u}, t) \equiv L(\mathbf{z}, \mathbf{w}, t)$), the equality reduces to the inner products of the costates and the vector fields:

$$\boldsymbol{\lambda}_x^\top \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t) = \boldsymbol{\lambda}_z^\top \mathbf{g}(\mathbf{z}(t), \mathbf{w}(t), t). \quad (26)$$

Substituting the dynamic transformation from Eq. (23) into Eq. (26) yields:

$$\boldsymbol{\lambda}_x^\top \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t) = \boldsymbol{\lambda}_z^\top [\mathbf{J}_\rho(\mathbf{x})] \mathbf{f}(\mathbf{x}(t), \mathbf{u}(t), t). \quad (27)$$

Since this relation must hold for all admissible optimal vector fields $\mathbf{f}$ at infinite states, the necessary and sufficient condition relating the two costate vectors at any instant t along the extremal trajectory is established by equating the covector coefficients:

$$\boldsymbol{\lambda}_x(t)^\top = \boldsymbol{\lambda}_z(t)^\top [\mathbf{J}_\rho(\mathbf{x}(t))] \quad (28)$$

By taking the transpose of Eq. (28), we obtain the exact mapping to recover the target costate $\boldsymbol{\lambda}_z$ from the source costate $\boldsymbol{\lambda}_x$:

$$\boldsymbol{\lambda}_z(t) = [\mathbf{J}_\rho(\mathbf{x}(t))]^{-\top} \boldsymbol{\lambda}_x(t). \quad (29)$$

This variational coordinate transformation is exact and holds for any pair of coordinate charts, as long as the mapping ρ is continuous, invertible, and differentiable. It provides a powerful tool to transition between different coordinate representations while preserving the optimality conditions of the original problem.

3.3 Unified Covector Mapping Framework

The core theoretical contribution of this work is the synthesis of the discrete-to-continuous primal/dual mapping (CMT) and the geometric coordinate transformation into a single, unified mathematical framework. In practical trajectory optimization, onboard algorithms must bridge a double non-commutation gap: transitioning from a discretized numerical method to a continuous method, while simultaneously shifting from a primary coordinate frame to an alternate one.

We resolve this by formally composing the two mapping operators. The CMT operator Γ_x^N maps the finite-dimensional discrete multipliers to a set of discrete temporal samples approximating the continuous-time covector in the original reference frame:

$$\Gamma_x^N : \{\boldsymbol{\Lambda}^x, \boldsymbol{\Gamma}^x, \boldsymbol{\Psi}^x\} \mapsto \{\tilde{\boldsymbol{\lambda}}_x, \tilde{\boldsymbol{\gamma}}_x, \tilde{\boldsymbol{\psi}}_x\}. \quad (30)$$

The coordinate mapping operator $\Pi_{x \rightarrow z}$ utilizes the inverse-transpose Jacobian to map the continuous covector into the target frame:

$$\Pi_{x \rightarrow z} : \tilde{\lambda}_x(t) \mapsto \tilde{\lambda}_z(t) = [\mathbf{J}_\rho(\mathbf{x}(t))]^{-\top} \tilde{\lambda}_x(t). \quad (31)$$

The Unified Covector Map, $\mathcal{M}_{x \rightarrow z}^N$, is defined as the mathematical composition of these operators:

$$\mathcal{M}_{x \rightarrow z}^N = \Pi_{x \rightarrow z} \circ \Gamma_x^N. \quad (32)$$

Because the intermediate variable $\lambda(t)$ exists in the continuous domain, the unified framework is inherently grid-independent. A solution generated on an arbitrary dense grid (e.g., via an indirect ODE solver) can be mapped, interpolated, and re-discretized onto a completely different sparse grid (e.g., LGR collocation nodes) to warm-start an alternative solver.

However, a powerful analytical simplification occurs when transitioning between two direct optimization problems that utilize identical temporal grids, segment distributions, and transcription methods. Under these strict conditions, the temporal nodes and quadrature weights of the source and target frames are perfectly identical ($w_k^x = w_k^z$). Substituting this equivalence into the variational mapping yields the discrete generalized map:

$$\left(\frac{\Lambda_k^z}{w_k^z} \right) = [\mathbf{J}_\rho(\mathbf{x}_k)]^{-\top} \left(\frac{\Lambda_k^x}{w_k^x} \right). \quad (33)$$

Because the identical weights cancel out entirely, the discrete multiplier mapping bypasses the continuous interpolation step and reduces to a purely algebraic rotation:

$$\Lambda_k^z = [\mathbf{J}_\rho(\mathbf{x}_k)]^{-\top} \Lambda_k^x. \quad (34)$$

This property demonstrates that under identical grids, the mapping operations commute. The coordinate transformation can be applied directly to the discrete KKT multipliers without first mapping them to the continuous domain. Therefore, an equivalent dual Unified Covector Map operator can be defined as:

$$\mathcal{H}_{x \rightarrow z}^N = \Gamma_z^N \circ \Pi_{x \rightarrow z}. \quad (35)$$

A schematic of the complete Unified Covector Mapping Framework, illustrating the three-dimensional extension across transcription methods and coordinate charts (e.g., Cartesian, MEE, and COE), is presented in Figure 2.

3.4 Hybrid Algorithm Methodology

The Unified Covector Mapping Framework provides a general mechanism for transferring optimal solutions, including both primal and dual variables, between heterogeneous optimal control formulations. Rather than coupling a specific direct and indirect solver, the framework enables the construction of hybrid optimization strategies in which different formulations can be exploited according to their respective numerical advantages. Such flexibility is particularly attractive for onboard autonomous guidance, where computational efficiency, robustness, and autonomous recovery from failed optimizations are critical requirements. In an onboard setting, different phases of a mission may benefit from different optimization formulations. Direct transcription methods provide robust convergence from poor initial guesses, whereas indirect methods offer rapid convergence and reduced computational cost once an accurate costate estimate is available. The proposed framework




Figure 2. Schematic of the Unified Covector Mapping Framework, which extend the original CMT diagram to a third dimension of coordinate transformation.

enables autonomous transitions between these formulations by consistently transferring optimality information across solver architectures and coordinate representations.

Within the context of the Unified Covector Mapping Framework a hybrid algorithm is proposed, exploiting both the mapping between direct and indirect methods and between MEE and Cartesian coordinates. It employs a state-of-the-art indirect shooting method formulated with exact switching detection,¹⁶ and the LGR pseudospectral collocation described in Section 2.3 with adaptive *hp*-mesh refinement¹⁷ and using IPOPT¹⁸ or SNOPT¹⁹ as NLP solvers.

Solving fuel-optimal problems via pseudospectral direct collocation is notoriously challenging, as discrete transcriptions relying on global polynomials struggle to accurately capture discontinuous bang-bang control structures without severe and computationally expensive mesh refinement. Conversely, these direct methods are highly efficient when solving smooth energy-optimal problems, particularly when coupled with MEE to handle multi-revolution trajectories. On the other hand, indirect methods suffer from extreme sensitivity to the initial guess; however, once provided with an energy-optimal solution, they can leverage a robust homotopy continuation strategy to systematically derive the discontinuous fuel-optimal solution. To ease convergence and increase robustness, the proposed hybrid pipeline extracts the optimal structure by exploiting the distinct advantages of both approaches while mitigating their respective weaknesses. The methodology operates through the following steps:

1. **Direct-MEE Energy Optimization:** The trajectory is first obtained by solving the energy-optimal problem through LGR pseudospectral transcription in MEE coordinates. The smooth control profile associated with the quadratic objective and the regularity of the MEE dynamics provide a large convergence domain while requiring only a relatively coarse adaptive mesh.
2. **Covector Mapping:** The discrete NLP multipliers are subsequently transformed into continuous Cartesian costates through the Unified Covector Mapping Framework. This operation simultaneously transfers the solution from the direct transcription to the indirect formulation

and from MEE to Cartesian coordinates, preserving the first-order optimality information required by the indirect solver.

3. **Indirect Warm-Start:** Because the mapped costates satisfy the analytical covector relationships derived in Section 3.1, the indirect solver is initialized with a physically consistent solution of the adjoint variables, substantially leading to immediate convergence of the energy-optimal problem in few iterations.
4. **Homotopy Continuation:** Once the smooth energy-optimal solution has converged, the indirect formulation provides a natural framework for homotopy continuation toward the discontinuous fuel-optimal solution. As the continuation parameter decreases, the switching structure emerges without requiring its manual identification.

Although this work focuses on a direct-to-indirect transition between MEE and Cartesian formulations, the proposed framework is not restricted to this particular pipeline. Any optimization architecture for which the corresponding coordinate Jacobians and transcription-specific covector mappings are available can be incorporated within the same framework. Consequently, additional hybrid strategies can be constructed by combining different state representations, transcription methods, or continuation techniques according to the numerical characteristics of the problem under consideration.

From an autonomous guidance perspective, this capability enables an optimization process to adapt its numerical representation during execution. For example, an onboard guidance system may employ a robust direct method during global trajectory initialization, transition to an indirect formulation for computationally efficient refinement, and subsequently return to alternative coordinate representations as mission conditions evolve. Because the transfer of optimality information is performed analytically, these transitions require no heuristic initialization of the adjoint variables, improving both robustness and computational efficiency.

4 NUMERICAL RESULTS

In this section, numerical results based on the developed theoretical framework are presented. In particular, the mapping of costates is verified with high accuracy both between direct and indirect methods and between Cartesian and MEE coordinate frames. The performance of the proposed hybrid algorithm is evaluated across a dataset of hundreds of interplanetary scenarios.

4.1 Direct-indirect costate mapping

Figures 3 and 4 show the results obtained with the PS direct collocation and indirect shooting for a fuel-optimal, low-thrust interplanetary transfer from Earth to Venus. The scenario assumes an initial spacecraft mass of 1500 kg, a maximum thrust of 0.33 N, a specific impulse of 3800 s, and a time of flight of 1000 days. For the PS method, a local collocation scheme with three segments and 40 nodes per segment was utilized. Both methods use Cartesian coordinates and are solved independently from each other.

As demonstrated in Table 1 and Figure 4, the mapping yields excellent costate estimation. This high-fidelity numerical agreement confirms that Lagrange multipliers found with LGR pseudospectral collocation are accurately mapped via the CMT to the original costates of the continuous-time problem. Small remaining differences are due to the numerical approximation of the discretization, and are not introduced by inaccuracies within the mapping framework itself.




Figure 3. Trajectory, control and mass of the fuel-optimal Earth-Venus transfer computed with direct and indirect methods.




Figure 4. Comparison of costates and mapped NLP Lagrange multipliers for the fuel-optimal Earth-Venus transfer.

Table 1. Comparison of initial costates $\lambda(t_0)$ obtained via indirect shooting and mapped direct Lagrange multipliers for the Earth-Venus transfer. The two sets of costates are obtained independently, without reciprocal warm-starting.

Costate	Indirect Shooting	Mapped Direct (CMT)	Absolute Error
λ_x	0.607623332345686	0.607389744942784	$2.3359e - 04$
λ_y	0.159410707754726	0.159513968260261	$1.0326e - 04$
λ_z	0.247186183801536	0.246006851456459	$1.1793e - 03$
λ_{vx}	-0.150794413483427	-0.150651726712936	$1.4269e - 04$
λ_{vy}	0.612951804656612	0.612997723535499	$4.5919e - 05$
λ_{vz}	-0.129310784978976	-0.130165208928353	$8.5442e - 04$
λ_m	0.138768747804402	0.138713168710153	$5.5579e - 05$

This excellent costate estimation has also been verified for both energy-optimal and time-optimal formulations. Consequently, the mapping of costates via the CMT can be effectively used to warm-start indirect methods using solutions from direct approaches, and vice versa. This enables the development of hybrid algorithms that leverage the strengths of both formulations while mitigating their respective limitations. In particular, solving this fuel-optimal problem via indirect shooting is highly sensitive to the initial guess, often resulting in convergence to different local optima across runs. In contrast, direct methods exhibit lower sensitivity to initial conditions; however, they require solving large-scale NLP problems involving multiple segments and a high number of collocation nodes to adequately capture the bang-bang control structure. Even with this increased resolution, the switching times are typically not identified with the same accuracy as those obtained from indirect methods.

4.2 Indirect-MEE to Direct-Cartesian costate mapping

Multiple-revolution trajectories are traditionally challenging to solve using Cartesian coordinates. Figure 5 illustrates a time-optimal GTO to GEO transfer for a 1500 kg spacecraft with a maximum thrust of 11 N. The problem was initially solved via indirect shooting using MEE. It is worth noting that indirect approaches can also exploit robust thrust continuation strategies to efficiently solve multiple-revolution transfers with decreasing thrust.²⁰ Utilizing the Unified Covector Mapping Framework, this MEE-based solution was subsequently mapped to initialize a direct collocation NLP in Cartesian coordinates.

The plotted solutions are perfectly superimposed, confirming the accuracy of the coordinate transformation. Furthermore, despite the high density of temporal nodes, the Cartesian direct solution converged in only a few iterations due to the highly accurate dual warm-start provided by the costates mapped to the discrete Lagrange multipliers. Note that the results are plotted in Cartesian coordinates for both approaches to facilitate a direct visual comparison.

This simulation demonstrates the validity of the costate mapping across both different optimization methods and coordinate frames, ultimately validating the core theoretical contribution of the Unified Covector Mapping Framework.




Figure 5. Trajectory and mapped costates of the GTO-GEO transfer computed with MEE-indirect and Cartesian-direct methods.

4.3 Direct-indirect hybrid algorithm results

The hybrid algorithm outlined in Section 3.4 was evaluated on 450 scenarios of multiple-revolution interplanetary transfers from Earth to Venus and Mars, developed for the ERC-funded *EXTREMA* project.²¹ Its performance was compared against a standard indirect shooting baseline. This baseline method was initialized using an ACT-based guess and exploited hyperbolic tangent smoothing alongside an energy-to-fuel homotopy.¹⁶ To account for the pseudo-random nature of the ACT guess, 10 independent runs were performed for each scenario using the indirect method.

The hybrid algorithm achieved successful convergence in 79.69% of cases, whereas the mean convergence rate of the baseline indirect method was 70.56%. Figure 6 illustrates the differences in normalized fuel mass and CPU time between the hybrid and indirect approaches for the commonly converged scenarios. (Note that the initial wet mass is normalized to one). In these plots, positive values (colored in red) indicate scenarios where the hybrid algorithm outperformed the indirect method by achieving lower fuel consumption or shorter computational times. Values for the indirect method represent the mean across the 10 independent runs.




Figure 6. Normalized fuel mass saved and CPU time saved for each of the scenarios converged with both methods.

On average, the hybrid algorithm demonstrates an absolute 9% increase in the convergence rate and a 46% reduction in CPU time. Furthermore, when evaluating the subset of scenarios where the two methods converged to different local optima, the hybrid approach yielded a 4.4% reduction in fuel consumption compared to the indirect shooting average.

These performance gains directly reflect the theoretical strengths of the unified framework. The 46% reduction in computational time occurs because the highly accurate mapped costates allow the indirect solver to immediately initialize near the true solution, bypassing the extensive exploratory iterations typically required when starting from a pseudo-random guess. Similarly, the fuel savings highlight the direct method’s ability to locate a superior basin of attraction, preventing the subsequent indirect refinement from getting trapped in suboptimal local minima. Finally, while the indirect method’s reliance on the ACT guess introduces inherent variance across runs, the hybrid approach provides a highly consistent pipeline. This combined robustness, efficiency, and predictability—enabled by seamlessly mapping solutions across different problem formulations—makes the unified hybrid approach a highly compelling candidate for onboard autonomous guidance systems.

5 CONCLUSION

This paper presents a Unified Covector Mapping Framework to address the challenges related to low-thrust trajectory optimization, specifically tailored for autonomous onboard spacecraft guidance. By mathematically composing the discrete-to-continuous Covector Mapping Theorem with continuous-time variational coordinate transformations, the framework establishes a rigorous bridge between direct and indirect optimization techniques and geometric reference frames. This theoretical foundation enables treating the dual variables as transportable mathematical assets, allowing trajectory optimization pipelines to exploit the distinct advantages of specific methods and coordinate representations at different stages of the solution process.

Leveraging this framework, a novel hybrid guidance algorithm is introduced to mitigate the high sensitivity of indirect methods and the computational burden of direct methods. The proposed strategy utilizes a computationally efficient direct energy-optimal solution formulated in MEE, which is subsequently mapped to warm-start an indirect Cartesian shooting method for accurate energy-to-fuel homotopy. Numerical validation across a dataset of 450 interplanetary transfer scenarios demonstrates the robust advantages of this approach. Compared to standard indirect methods relying on traditional initial guess generation, the unified hybrid algorithm achieves a 9% increase in convergence reliability, a 4.4% reduction in propellant mass, and a 46% decrease in computational time. Ultimately, the proposed unified mapping methodology paves the way for innovative guidance strategies with increased robustness and efficiency in solving non-smooth optimal control problems. Future work will focus on expanding the hybrid pipeline to dynamically bootstrap complex environmental and operational constraints, such as eclipses or thrust pointing limits, and deploying the algorithm for processor-in-the-loop and hardware-in-the-loop testing to definitively validate its viability for next-generation self-driving spacecraft.

Acknowledgments

This research is part of EXTREMA, a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant Agreement No. 864697).

REFERENCES

- [1] A. E. Bryson, *Applied optimal control: optimization, estimation and control*. Routledge, 2018.
- [2] D. G. Hull, *Optimal control theory for applications*. Springer Science & Business Media, 2013.
- [3] G. Di Domenico, E. Andreis, A. Carlo Morelli, G. Merisio, V. Franzese, C. Giordano, A. Morselli, P. Panicucci, F. Ferrari, and F. Topputo, “The ERC-funded EXTREMA project: Achieving self-driving interplanetary cubesats,” *Modeling and Optimization in Space Engineering: New Concepts and Approaches*, pp. 167–199, Springer, 2022.
- [4] A. Mannocchi, D. Perico, G. Oliveira Pinho, G. Di Domenico, A. Morselli, F. Topputo, *et al.*, “An Overview of EXTREMA Project Pillar II: Validating Autonomous Guidance Algorithms via HIL Testing,” *IAF Space Systems Symposium (75th International Astronautical Congress)*, International Astronautical Federation (IAF), 2024, pp. 1157–1167.
- [5] D. Morante, M. Sanjurjo Rivo, and M. Soler, “A survey on low-thrust trajectory optimization approaches,” *Aerospace*, Vol. 8, No. 3, 2021, p. 88.
- [6] F. Topputo and C. Zhang, “Survey of Direct Transcription for Low-Thrust Space Trajectory Optimization with Applications,” *Abstract and Applied Analysis*, Vol. 2014, Wiley Online Library, 2014, p. 851720.
- [7] P. J. Enright and B. A. Conway, “Discrete approximations to optimal trajectories using direct transcription and nonlinear programming,” *Journal of Guidance, Control, and Dynamics*, Vol. 15, No. 4, 1992, pp. 994–1002.
- [8] J. L. Junkins and E. Taheri, “Exploration of alternative state vector choices for low-thrust trajectory optimization,” *Journal of Guidance, Control, and Dynamics*, Vol. 42, No. 1, 2019, pp. 47–64.
- [9] C. Hofmann, A. C. Morelli, F. Topputo, *et al.*, “Impact of different coordinate sets on the performance of convex low-thrust trajectory optimization,” *33rd AAS/AIAA space flight mechanics meeting*, 2023, pp. 1–20.
- [10] F. Spada, M. Sagliano, and F. Topputo, “Direct–indirect hybrid strategy for optimal powered descent and landing,” *Journal of Spacecraft and Rockets*, Vol. 60, No. 6, 2023, pp. 1787–1804.
- [11] N. Michelotti, C. Giordano, F. Topputo, *et al.*, “Low-thrust trajectory optimization via Covector Mapping Principle: bridging direct and indirect methods for spacecraft autonomous guidance,” *76th International Astronautical Congress (IAC 2025)*, 2025, pp. 1–11.
- [12] I. M. Ross, “A historical introduction to the convector mapping principle,” *Proceedings of Astrodynamics Specialists Conference*, Naval Postgraduate School (US), 2005.
- [13] E. Taheri, V. Arya, and J. L. Junkins, “Costate mapping for indirect trajectory optimization,” *Astrodynamics*, Vol. 5, No. 4, 2021, pp. 359–371.
- [14] Y. Wang and F. Topputo, “Indirect optimization of power-limited asteroid rendezvous trajectories,” *Journal of Guidance, Control, and Dynamics*, Vol. 45, No. 5, 2022, pp. 962–971.
- [15] D. Garg, M. Patterson, W. W. Hager, A. V. Rao, D. A. Benson, and G. T. Huntington, “A unified framework for the numerical solution of optimal control problems using pseudospectral methods,” *Automatica*, Vol. 46, No. 11, 2010, pp. 1843–1851.
- [16] C. Zhang, F. Topputo, F. Bernelli-Zazzera, and Y.-S. Zhao, “Low-thrust minimum-fuel optimization in the circular restricted three-body problem,” *Journal of Guidance, Control, and Dynamics*, Vol. 38, No. 8, 2015, pp. 1501–1510.
- [17] M. A. Patterson, W. W. Hager, and A. V. Rao, “A ph mesh refinement method for optimal control,” *Optimal Control Applications and Methods*, Vol. 36, No. 4, 2015, pp. 398–421.
- [18] A. Wächter and L. T. Biegler, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Mathematical programming*, Vol. 106, 2006, pp. 25–57.
- [19] P. E. Gill, W. Murray, and M. A. Saunders, “SNOPT: An SQP algorithm for large-scale constrained optimization,” *SIAM review*, Vol. 47, No. 1, 2005, pp. 99–131.
- [20] Y. Wang, X. Hou, and F. Topputo, “Thrust continuation of time-optimal orbital transfers with soft terminal conditions,” *Astrophysics and Space Science*, Vol. 369, No. 4, 2024, p. 39.
- [21] N. Michelotti, G. Oliveira Pinho, and F. Topputo, “EXTREMA: SCP-Optimized Low-Thrust Interplanetary Trajectories,” Feb. 2026, 10.5281/zenodo.18485427.