Physics-Informed ADMM for Path-Constrained Continuous-Time Optimal Control

Eunbyeol Cho[1], Dohoon Lee[1] and Jong-Han Kim[1†]

Abstract— **This paper proposes a physics-informed alternating direction method of multipliers (PI-ADMM) framework for constrained continuous-time optimal control and applies it to a 3-DoF powered descent guidance (PDG) problem. The method combines an indirect primal update based on Pontryagin's Minimum Principle, computed from the augmented Hamiltonian system, with an outer-loop ADMM scheme for path-constraint enforcement. The primal trajectories are parameterized using Chebyshev-based constrained expressions and computed by minimizing the augmented PMP residuals. An optimality analysis shows that converged fixed points of the proposed continuous-time ADMM iterations satisfy the first-order necessary conditions under the convex-set assumption. Numerical results show solution quality comparable to IPOPT and lossless convexification in terms of fuel consumption, constraint satisfaction, and propagation error, while achieving the shortest solve time among the tested solvers.**

I. Introduction

Path-constrained continuous-time optimal control problems arise in many applications, for example powered descent guidance (PDG), where one seeks dynamically feasible and constraint-satisfying trajectories under limited computation. In practical settings, the solver may need to operate in an online replanning loop while handling nonlinear dynamics and state/control constraints. This motivates computationally efficient methods that preserve continuous-time optimality structure while enforcing path feasibility effectively.

Existing approaches are broadly classified as direct and indirect. Direct methods transcribe the continuous-time problem into a nonlinear or convex program, making path constraints relatively straightforward to impose at discretization nodes [1]–[5], but often at the cost of large-scale transcription, successive convexification, or heavy numerical solvers. Indirect methods instead apply Pontryagin's Minimum Principle (PMP) to derive a two-point boundary value problem (TPBVP) [6]; they preserve the continuous-time optimality structure and can yield highly accurate solutions, but are often sensitive to the initial costate estimate and less convenient for complicated path constraints. This challenge has motivated recent interest in physics-informed indirect solvers.

Recently, physics-informed neural approaches [7] have been explored for solving the TPBVPs arising from indirect optimal control. In this line of research, the Theory of Functional Connections (TFC) provides constrained expressions that analytically satisfy boundary conditions, whereas its neural extension, the Extreme Theory of Functional Connections (X-TFC), employs a single-layer neural network trained via the Extreme Learning Machine (ELM) algorithm to obtain computationally efficient continuous-time approximations [8]–[10]. These ideas have been successfully applied to several optimal control problems [11]–[14]. However, earlier studies in this research line either focused on unconstrained problems, handled bounded controls through saturation and continuation, or did not provide a general treatment of state, control, and mixed state-control inequality path constraints. A recent extension, Physics-Informed Pontryagin Neural Networks for Path-Constrained Optimal Control Problems [15], addresses this gap by introducing additional Lagrange multipliers directly into the Hamiltonian and enforcing complementarity through the Fischer–Burmeister function. While this represents an important step toward constrained indirect neural optimal control, it also requires optimizing the path multipliers and adding an additional complementarity residual term to the loss; moreover, because the Fischer–Burmeister function is nonsmooth at the origin, the method may increase tuning burden and can introduce numerical difficulties near active-set transitions.

This paper addresses these limitations by combining an indirect TPBVP solver with an outer-loop alternating direction method of multipliers (ADMM) framework. Specifically, the primal update preserves the continuous-time PMP structure through an augmented Hamiltonian formulation, while path feasibility is enforced through projection-based ADMM updates. Based on this idea, we propose a *Physics-Informed ADMM* (PI-ADMM) framework for constrained continuous-time optimal control. The primal TPBVP is solved using a physics-informed neural architecture with Chebyshev-based constrained expressions [12], [16], which analytically satisfy boundary and transversality conditions, while the outer loop updates auxiliary and dual variables through operator splitting. In this way, the proposed method combines the continuous-time optimality structure of indirect methods with the practical constraint-handling flexibility of first-order splitting methods.

The main contributions of this paper are summarized as follows:

This work was in part supported by the Korea Research Institute for Defense Technology Planning and Advancement (KRIT) under Grant KRIT-CT-22-030 through the Reusable Unmanned Space Vehicle Research Center, funded by the Korea government (Defense Acquisition Program Administration, DAPA), and in part by the Space Challenge Program funded by Korea AeroSpace Administration (KASA), under Grant RS-2025-16063807. Both funding sources contributed equally to this work. *(Corresponding author: Jong-Han Kim.)*

[1]The authors are with the Department of Aerospace Engineering and the Program in Aerospace Systems Convergence, Inha University, Incheon, Republic of Korea. eunbyeeol@inha.edu, dohoonlee@inha.edu, jonghank@inha.ac.kr

- We propose a continuous-time ADMM reformulation for path-constrained optimal control that decouples path-feasibility enforcement from the indirect primal TPBVP solve via variable splitting and projection.
- We show that, under the convex-set assumption, converged fixed points of the proposed iterations satisfy the first-order necessary conditions of the original continuous-time problem.
- We validate the proposed PI-ADMM framework on a 3-DoF PDG problem and show that it achieves solution quality comparable to IPOPT and LCvx while providing the shortest solve time among the tested methods.

II. Preliminaries

A. Continuous-Time Optimal Control Problem

We consider a continuous-time nonlinear optimal control problem with path constraints. Over $t \in [t_0, t_f]$, let $x(t) \in \mathbb{R}^{n_x}$ and $u(t) \in \mathbb{R}^{n_u}$ denote the state and control, respectively. The constrained OCP is given by (1), where Φ and $\mathcal{L}$ denote the terminal and running costs, f is the nonlinear dynamics, and $g(t,x,u) \in \mathcal{C}$ enforces path feasibility over a closed convex set $\mathcal{C}$.

$$\begin{aligned} \underset{x,u}{\text{minimize}} \quad & \Phi(t_f, x(t_f)) + \int_{t_0}^{t_f} \mathcal{L}(t, x(t), u(t))dt \\ \text{subject to} \quad & \dot{x}(t) = f(t, x(t), u(t)), \\ & g(t, x(t), u(t)) \in \mathcal{C}, \\ & x(t_0) = x_0, \end{aligned} \tag{1}$$

To derive first-order necessary conditions, we introduce the Hamiltonian with costate $\lambda(t)$ and path multiplier $\mu(t)$:

$$H(t,x,u,\lambda,\mu) = \mathcal{L}(t,x,u) + \lambda^\top f(t,x,u) + \mu^\top g(t,x,u),$$

where the following PMP conditions are necessary for local optimality:

$$\begin{aligned} \dot{x} &= \tfrac{\partial H}{\partial \lambda} = f(t,x,u), \\ -\dot{\lambda} &= \tfrac{\partial H}{\partial x} = \tfrac{\partial \mathcal{L}}{\partial x} + (\tfrac{\partial f}{\partial x})^\top \lambda + (\tfrac{\partial g}{\partial x})^\top \mu, \\ 0 &= \tfrac{\partial H}{\partial u} = \tfrac{\partial \mathcal{L}}{\partial u} + (\tfrac{\partial f}{\partial u})^\top \lambda + (\tfrac{\partial g}{\partial u})^\top \mu. \end{aligned} \tag{2}$$

In addition, primal and dual feasibility are characterized by

$$g(t,x,u) \in \mathcal{C}, \qquad \mu(t) \in \mathcal{N}_{\mathcal{C}}(g(t,x,u)), \tag{3}$$

where $\mathcal{N}_{\mathcal{C}}(\cdot)$ denotes the normal cone to the set $\mathcal{C}$. If $\mathcal{C} = \mathbb{R}^m_-$, this normal-cone condition is equivalent to the componentwise KKT conditions $g(t,x,u) \le 0$, $\mu(t) \ge 0$, and $\mu(t) \odot g(t,x,u) = 0$. Together with boundary conditions, these conditions form a TPBVP. However, coupled nonlinear path constraints make shooting-based indirect methods highly sensitive to the initial costate guess, motivating the constraint-handling strategy developed in the subsequent

B. Alternating Direction Method of Multipliers

First-order methods handle constraints via gradient-based updates and proximal operations. Among them, ADMM is particularly effective when constraints can be isolated through auxiliary variables and projection steps.

Consider the constrained problem $\text{minimize}_{x \in \mathcal{C}}\, \phi(x)$. Introducing an auxiliary variable z yields the equivalent form (4), where $I_{\mathcal{C}}(z)$ is the indicator function of $\mathcal{C}$ (0 if $z \in \mathcal{C}$, $+\infty$ otherwise).

$$\begin{aligned} \text{minimize}_x \quad & \phi(x) + I_{\mathcal{C}}(z) \\ \text{subject to} \quad & x = z \end{aligned} \tag{4}$$

The augmented Lagrangian with the penalty parameter $\rho > 0$ is given by:

$$\mathcal{L}_\rho(x,z,\mu) = \phi(x) + I_{\mathcal{C}}(z) + \mu^\top(x - z) + \tfrac{\rho}{2}\|x - z\|^2.$$

Using the scaled dual variable $y := \mu/\rho$, we obtain the following scaled form:

$$\mathcal{L}_\rho(x,z,y) = \phi(x) + I_{\mathcal{C}}(z) + \tfrac{\rho}{2}\|x - z + y\|^2 - \tfrac{\rho}{2}\|y\|^2 \tag{5}$$

ADMM then proceeds by alternating minimization and dual ascent, leading to iterative updates:

$$\begin{aligned} x^{k+1} &= \arg\min\nolimits_x \ \phi(x) + \tfrac{\rho}{2}\|x - z^k + y^k\|^2, \\ z^{k+1} &= \Pi_{\mathcal{C}}(x^{k+1} + y^k), \\ y^{k+1} &= y^k + x^{k+1} - z^{k+1}. \end{aligned} \tag{6}$$

In particular, the z-update reduces to the orthogonal projection $\Pi_{\mathcal{C}}(\cdot)$ onto the feasible set. This splitting-and-projection structure is leveraged in this paper to handle coupled path constraints in nonlinear optimal control problems.

III. Proposed Method: PI-ADMM

Our goal is to solve the continuous-time constrained OCP (1) using an indirect TPBVP approach while handling path constraints via an outer-loop splitting scheme.

A. Continuous-Time ADMM Framework

We embed ADMM into the continuous-time formulation by introducing an auxiliary variable $z(t)$ for the path constraint and enforcing the coupling constraint $g(t,x,u) - z(t) = 0$. This yields the equivalent augmented problem:

$$\begin{aligned} \underset{x,u,z}{\text{minimize}} \quad & \Phi(t_f, x(t_f)) + \int_{t_0}^{t_f} \left[\mathcal{L}(t,x,u) + I_{\mathcal{C}}(z(t))\right] dt \\ \text{subject to} \quad & \dot{x} = f(t,x,u), \\ & g(t,x,u) - z(t) = 0, \\ & x(t_0) = x_0, \end{aligned} \tag{7}$$

where feasibility with respect to $\mathcal{C}$ is encoded through the indicator function $I_{\mathcal{C}}(\cdot)$. Following the scaled-form ADMM in (5)-(6), we define the augmented cost functional (8) with the quadratic penalty term $\tilde{\mathcal{L}}(t,x,u,y,z)$ in (9), where $y(t)$ denotes the scaled dual variable:

$$\mathcal{J} = \Phi\big(t_f, x(t_f)\big) + \int_{t_0}^{t_f} \big[\tilde{\mathcal{L}}(t,x,u,y,z) + I_{\mathcal{C}}(z(t))\big] dt, \tag{8}$$

$$\tilde{\mathcal{L}}(t,x,u,y,z) = \mathcal{L}(t,x,u) + \tfrac{\rho}{2}\|g(t,x,u) - z + y\|^2. \tag{9}$$

The resulting continuous-time ADMM iterations consist of the following three steps. For fixed (z^k, y^k), the primal update solves a nonlinear OCP with dynamics (1) and the augmented running cost (9):

$$(x^{k+1}, u^{k+1}) = \arg\min_{x,u} \ \mathcal{J}_k(x,u) \\ \text{s.t. } \dot{x} = f(t,x,u), \ x(t_0) = x_0, \tag{10}$$

where $\mathcal{J}_k(x,u) := \Phi\big(t_f, x(t_f)\big) + \int_{t_0}^{t_f} \tilde{\mathcal{L}}\big(t,x,u,y^k,z^k\big)dt$. Next, the auxiliary variable is updated by projection onto $\mathcal{C}$,

$$z^{k+1} = \Pi_{\mathcal{C}}\big(g(t,x^{k+1},u^{k+1}) + y^k\big), \tag{11}$$

and the scaled dual variable is updated via a dual-ascent step enforcing $g = z$,

$$y^{k+1} = y^k + g(t,x^{k+1},u^{k+1}) - z^{k+1}. \tag{12}$$

B. Primal Update via PINN

1) PINN for TPBVP: In (10), the primal update is a continuous-time nonlinear optimal control subproblem. We treat this from an indirect optimal control perspective. Specifically, using the augmented Lagrangian in (9), we define the augmented Hamiltonian:

$$\tilde{H}(t,x,u,\lambda;y,z) = \tilde{\mathcal{L}}(t,x,u,y,z) + \lambda^\top f(t,x,u), \tag{13}$$

where $\lambda(t)$ denotes the costate associated with the system dynamics.

Applying PMP to (13) yields the first-order necessary conditions in (14), consisting of the state equation, the costate equation, and the stationarity condition with respect to the control.

$$\begin{aligned} \dot{x} &= \tfrac{\partial \tilde{H}}{\partial \lambda} = f(t,x,u) \\ -\dot{\lambda} &= \tfrac{\partial \tilde{H}}{\partial x} = \tfrac{\partial \mathcal{L}}{\partial x} + (\tfrac{\partial f}{\partial x})^\top \lambda + \rho(\tfrac{\partial g}{\partial x})^\top (g - z + y) \\ 0 &= \tfrac{\partial \tilde{H}}{\partial u} = \tfrac{\partial \mathcal{L}}{\partial u} + (\tfrac{\partial f}{\partial u})^\top \lambda + \rho(\tfrac{\partial g}{\partial u})^\top (g - z + y) \end{aligned} \tag{14}$$

Together with the boundary conditions, (14) defines a TPBVP for the primal variables $(x(t), \lambda(t), u(t))$. We solve this TPBVP using a PINN by approximating (x, λ, u) with neural function approximators and minimizing the residuals of (14) over collocation points on $[t_0, t_f]$.

2) Chebyshev NN with Constrained Expression: We parameterize the trajectories using a single-layer Chebyshev neural network with first-kind Chebyshev polynomials as hidden features:

$$T_0(x) = 1, \qquad T_1(x) = x, \\ T_{n+1}(x) = 2xT_n(x) - T_{n-1}(x). \tag{15}$$

The time domain $t \in [t_0, t_f]$ is linearly mapped to the standard interval $\tau \in [-1, 1]$ via

$$\tau(t) = 2\tfrac{t-t_0}{t_f-t_0} - 1, \quad \tfrac{d\tau}{dt} = \tfrac{2}{t_f-t_0}. \tag{16}$$

Using $T_i(\tau)$ as basis functions, an unconstrained free function $\psi(t)$ is expanded as:

$$\psi(t) = \textstyle\sum_{i=0}^{L-1} \xi_i T_i\big(\tau(t)\big), \tag{17}$$

where L is the number of basis functions and ξ denotes the trainable coefficients.

We use independent free-function parameterizations $\psi_x(t)$, $\psi_\lambda(t)$, and $\psi_u(t)$ for $x(t)$, $\lambda(t)$, and $u(t)$, with coefficient vectors $\xi_x \in \mathbb{R}^{n_x L}$, $\xi_\lambda \in \mathbb{R}^{n_\lambda L}$, and $\xi_u \in \mathbb{R}^{n_u L}$. To analytically enforce boundary conditions, we adopt the constrained expression (CE) framework of the Theory of Functional Connections (TFC) [8]. For example, when the state satisfies the endpoint constraints $x(t_0) = x_0$ and $x(t_f) = x_{\text{des}}$, we construct $x(t)$ as the constrained expression (18).

$$x(t) = \psi_x(t) + \tfrac{t_f - t}{t_f - t_0}(x_0 - \psi_x(t_0)) + \tfrac{t - t_0}{t_f - t_0}(x_{\text{des}} - \psi_x(t_f)) \tag{18}$$

This expression satisfies both boundary conditions for any choice of $\psi_x(\cdot)$. Differentiating (18) with respect to time yields (19).

$$\dot{x}(t) = \dot{\psi}_x(t) - \tfrac{1}{t_f - t_0}(x_0 - \psi_x(t_0)) + \tfrac{1}{t_f - t_0}(x_{\text{des}} - \psi_x(t_f)) \tag{19}$$

The derivative of the free function is obtained using the chain rule:

$$\dot{\psi}(t) = \dot{\tau}(t) \textstyle\sum_{i=0}^{L-1} \xi_i \tfrac{dT_i}{d\tau}\big(\tau(t)\big) = \tfrac{2}{t_f - t_0} \textstyle\sum_{i=0}^{L-1} \xi_i T_i'(\tau(t)). \tag{20}$$

The same construction is applied to the costate, where the CE is formed to satisfy the transversality conditions implied by the terminal cost and terminal constraints.

3) Residual Collocation and Least-Squares Update: We evaluate the TPBVP conditions in (14) at a set of collocation points $\{t_i\}_{i=1}^N \subset [t_0, t_f]$ and define the pointwise PMP residual at t_i as follows:

$$r_i(\xi_x, \xi_\lambda, \xi_u; y^k, z^k) := \begin{bmatrix} \dot{x}(t_i) - f(t_i, x, u) \\ \dot{\lambda}(t_i) + \frac{\partial \tilde{H}}{\partial x}(t_i, x, u, \lambda) \\ \frac{\partial \tilde{H}}{\partial u}(t_i, x, u, \lambda) \end{bmatrix} \tag{21}$$

Stacking these residuals yields $r(\xi_x, \xi_\lambda, \xi_u; y^k, z^k)$, and the primal update is formulated as the nonlinear least-squares problem:

$$\text{minimize}_{\xi_x, \xi_\lambda, \xi_u} \ \|r(\xi_x, \xi_\lambda, \xi_u; y^k, z^k)\|^2. \tag{22}$$

We solve (22) using an iterative least-squares procedure in the inner loop. Let $J(\xi) := \partial r / \partial \xi$ denote the Jacobian of the residual with respect to the concatenated coefficient vector $\xi := \left[\xi_x^\top, \xi_\lambda^\top, \xi_u^\top\right]^\top \in \mathbb{R}^p$. Since the residual is evaluated at N collocation nodes with d_r components per node, the Jacobian has the structure $J(\xi) \in \mathbb{R}^{Nd_r \times p}$. At each inner iteration, we compute the Levenberg–Marquardt step

$$\big(J^\top J + \epsilon I\big)\Delta\xi = -J^\top r, \quad \xi \leftarrow \xi + \Delta\xi. \tag{23}$$

Here, $\epsilon > 0$ is a damping parameter that improves the numerical stability. In our implementation, $J(\xi)$ is obtained efficiently via automatic differentiation.

The PI-ADMM schematic is summarized in Fig. 1.

C. Optimality Analysis

This subsection presents an ideal fixed-point consistency result for the proposed PI-ADMM architecture. Specifically, the result assumes that the primal subproblem is solved to stationarity at each ADMM iteration. In Section IV-B, we instead use a practical inexact variant with only one inner step per outer iteration.

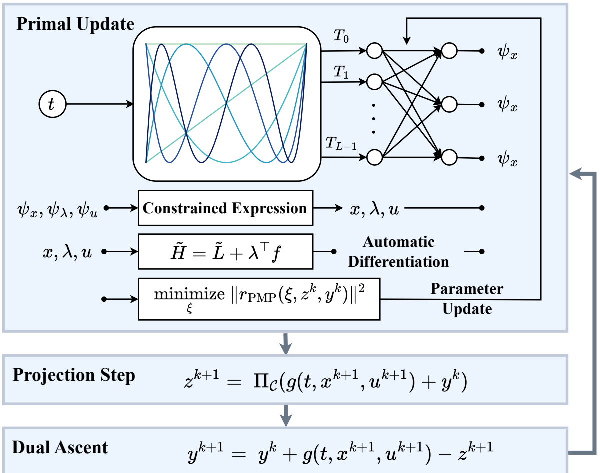

Fig. 1. Schematic of the proposed PI-ADMM framework for path-constrained continuous-time optimal control.

Proposition 1 (Fixed-Point Consistency)**.** *Assume that $\mathcal{C}$ is closed and convex. Suppose that the continuous-time ADMM iterations* (10)–(12)*, with the primal update solved to stationarity through* (14)*, converge to a fixed point $(x^\star, u^\star, \lambda^\star, z^\star, y^\star)$. Define the path multiplier by $\mu^\star := \rho y^\star$. Then $(x^\star, u^\star, \lambda^\star, \mu^\star)$ satisfies: (i) $g(t, x^\star(t), u^\star(t)) \in \mathcal{C}$ (primal feasibility), (ii) $\mu^\star(t) \in \mathcal{N}_{\mathcal{C}}(g(t, x^\star(t), u^\star(t)))$ (dual feasibility), and (iii) the PMP conditions of* (2)*. If $\mathcal{C} = \mathbb{R}^m_-$, the complementarity slackness holds almost everywhere in t.*

Proof: The z-update $z^{k+1} = \arg\min_z I_{\mathcal{C}}(z) + \frac{\rho}{2}\|z - (g(t, x^{k+1}, u^{k+1}) + y^k)\|^2$ is convex in z, hence its first-order optimality condition is the subdifferential inclusion:

$$0 \in \partial I_{\mathcal{C}}(z^{k+1}) + \rho\big(z^{k+1} - (g(t, x^{k+1}, u^{k+1}) + y^k)\big). \quad (24)$$

Since $\mathcal{C}$ is closed and convex, $\partial I_{\mathcal{C}}(z) = \mathcal{N}_{\mathcal{C}}(z)$, and (24) is equivalent to:

$$\rho\big(g(t, x^{k+1}, u^{k+1}) - z^{k+1} + y^k\big) \in \mathcal{N}_{\mathcal{C}}(z^{k+1}). \quad (25)$$

At a fixed point, the y-update (12) implies $g^\star(t) = z^\star(t)$ for all t, and therefore $g^\star(t) \in \mathcal{C}$, proving primal feasibility. Substituting $g^\star = z^\star$ into (25) yields $\rho\, y^\star(t) \in \mathcal{N}_{\mathcal{C}}(g^\star(t))$, hence $\mu^\star(t) \in \mathcal{N}_{\mathcal{C}}(g^\star(t))$, proving dual feasibility. If $\mathcal{C} = \mathbb{R}^m_-$, the normal-cone condition implies the componentwise complementarity slackness almost everywhere in time.

Finally, evaluating the augmented PMP conditions (14) at the fixed point and using $g^\star - z^\star = 0$, the augmented constraint terms in both the costate and stationary conditions reduce to the original PMP multiplier terms $(\partial g/\partial x)^\top \mu^\star$ and $(\partial g/\partial u)^\top \mu^\star$, respectively. Hence, the fixed point satisfies the state equation, costate equation, and stationarity condition in (2), which coincide with the PMP conditions of the original constrained OCP, completing the proof. □

Proposition 1 assumes an exact stationary primal update. The implemented solver instead uses an inexact primal step, whose effect is characterized below in terms of the augmented PMP residual and the ADMM primal/dual residuals.

Corollary 2 (Inexact primal-step consistency)**.** *Assume that $\mathcal{C}$ is closed and convex. Suppose that the primal step is computed inexactly so that*

$$\|r(\xi^{k+1}; y^k, z^k)\| \le \varepsilon_k, \quad (26)$$

where $r(\cdot; y^k, z^k)$ is the augmented PMP residual defined in (21)*. Let $g^{k+1} := g(t, x^{k+1}, u^{k+1})$, $r_p^{k+1} := g^{k+1} - z^{k+1}$, $r_d^{k+1} := z^{k+1} - z^k$, and define $\mu^{k+1} := \rho y^{k+1}$. If the iterate sequence is bounded and $\partial g/\partial x$ and $\partial g/\partial u$ remain bounded along it, then there exists a constant $c_g > 0$ such that*

$$\|e_{\mathrm{PMP}}^{k+1}\| \le \varepsilon_k + c_g \|r_d^{k+1}\|, \quad (27)$$

where e_{PMP}^{k+1} denotes the residual of the original PMP system (2) *evaluated at $(x^{k+1}, u^{k+1}, \lambda^{k+1}, \mu^{k+1})$. Moreover, $\mu^{k+1}(t) \in \mathcal{N}_{\mathcal{C}}(z^{k+1}(t))$ almost everywhere. Consequently, if $\varepsilon_k \to 0$, $\|r_p^k\| \to 0$, and $\|r_d^k\| \to 0$, then every cluster point of the inexact PI-ADMM iterates satisfies the conclusions of Proposition 1.*

Proof: (Sketch). Using (12), $g^{k+1} - z^k + y^k = y^{k+1} + r_d^{k+1}$. Hence the augmented PMP residual equals the original PMP residual, evaluated with $\mu^{k+1} = \rho y^{k+1}$, up to perturbation terms of order $\|r_d^{k+1}\|$; bounded $\partial g/\partial x$ and $\partial g/\partial u$ then yield (27). The projection step implies $y^{k+1} \in \mathcal{N}_{\mathcal{C}}(z^{k+1})$, and thus $\mu^{k+1} \in \mathcal{N}_{\mathcal{C}}(z^{k+1})$. Finally, $\|r_p^k\| \to 0$ gives $g^\star = z^\star \in \mathcal{C}$ at any cluster point, and $\varepsilon_k \to 0$, $\|r_d^k\| \to 0$ imply vanishing original PMP residual. Hence the limit satisfies Proposition 1. □

Remark 3 (Scope of the Convex-Set Assumption)**.** *The identity $\partial I_{\mathcal{C}} = \mathcal{N}_{\mathcal{C}}$ requires $\mathcal{C}$ to be closed and convex. In practice, the proposed framework can also be applied to non-convex feasible sets through projection operators that may be expansive. While nonexpansivity is a standard sufficient condition for convergence, expansive projections have also shown stable practical performance in first-order methods [17]–[20].*

IV. Numerical Experiments

A. Problem Definition

We validate the proposed PI-ADMM framework on a 3-DoF PDG problem [5], [17]:

$$\begin{aligned}
&\underset{u(\cdot),x(\cdot)}{\text{minimize}} && -m(t_f) + w_{\text{slew}} \int_{t_0}^{t_f} \|u(t)\|^2 dt && (28a)\\
&\text{subject to} && \dot{r} = v, \qquad \dot{v} = T/m + g_{\text{grav}}, && (28b)\\
& && \dot{m} = -\alpha\|T\|, \qquad \dot{T} = u, && (28c)\\
& && \|v\| \le v_{\max}, \qquad \|u\| \le \dot{T}_{\max}, && (28d)\\
& && m_{\text{dry}} \le m \le m_0, && (28e)\\
& && T_u \ge \cos(\gamma_{\text{p}})\|T\|, && (28f)\\
& && r_u \ge \tan(\gamma_{\text{gs}})\|r_{e:n}\|, && (28g)\\
& && T_{\min} \le \|T\| \le T_{\max}, && (28h)\\
& && r(t_0) = r_0,\ v(t_0) = v_0,\ m(t_0) = m_0, && (28i)\\
& && r(t_f) = r_{\text{des}},\ v(t_f) = v_{\text{des}}. && (28j)
\end{aligned}$$

TABLE I

NOMINAL PROBLEM DATA FOR THE 3-DOF PDG BENCHMARK USED IN THE NUMERICAL STUDY.

Parameter	**Value**	**Parameter**	**Value**
t_f	35 s	K	50
g_{grav}	9.81 m/s^2	I_{sp}	270 s
$T_{\min}$	40 kN	$T_{\max}$	60 kN
$\dot{T}_{\max}$	$0.05\,T_{\max}$	$v_{\max}$	35 m/s
γ_p	$10°$	γ_{gs}	$30°$
m_0	5000 kg	m_{dry}	3000 kg
r_0	$[800,\ 250,\ 100]^\top$ m	r_{des}	$[10,\ 0,\ 0]^\top$ m
v_0	$[-8,\ 10,\ 0]^\top$ m/s	v_{des}	$[0,\ 0,\ 0]^\top$ m/s

To explicitly model thrust slew rate limits, we augment the translational dynamics with the thrust vector as a state and treat the slew as the control input. The state and control are defined as $x = [r^\top, v^\top, m, T^\top]^\top \in \mathbb{R}^{10}$ and $u = \dot{T} \in \mathbb{R}^3$, where $r \in \mathbb{R}^3$ is position, $v \in \mathbb{R}^3$ is velocity, $m \in \mathbb{R}$ is mass, and $T \in \mathbb{R}^3$ is the thrust vector. The terminal cost maximizes the terminal mass, while the running cost penalizes the thrust slew, with w_{slew} denoting the associated weighting parameter. The dynamics (28b)–(28c) describe translational motion under thrust and gravity, propellant depletion, and the slew dynamics.

The path constraints enforce velocity, slew-rate, and dry-mass limits (28d)–(28e), together with thrust pointing and glide-slope constraints (28f)–(28g). The thrust-magnitude constraint (28h) produces a feasible nonconvex annulus set, which can induce an expansive projection (Remark 3). The boundary conditions are given in (28i)–(28j). Since the initial and terminal thrust is free and Φ depends on $m(t_f)$, the corresponding transversality conditions are:

$$\lambda_T(t_0) = 0, \qquad \lambda_T(t_f) = 0, \qquad \lambda_m(t_f) = -1. \tag{29}$$

These terminal conditions are incorporated when constructing the constrained expressions for the costates in the primal TPBVP solver.

B. Simulation Results

The simulation setup is summarized in Table I. Here, K denotes the number of time nodes, and α in the mass dynamics (28c) is the propellant mass flow coefficient, given by $\alpha = 1/(I_{\text{sp}} g_{\text{grav}})$. The proposed PI-ADMM solver is compared with two baselines: a nonlinear programming (NLP) approach implemented in CasADi and solved by IPOPT, and a lossless convexification (LCvx) formulation implemented in CVXPY and solved by MOSEK [2], [5]. These baselines represent, respectively, a generic nonlinear direct method and a convexified solution framework for the same class of PDG problems.

Fig. 2 shows the residual histories of the proposed PI-ADMM solver. In the numerical experiments, we use the inexact PI-ADMM variant of Corollary 2, with one inner Levenberg–Marquardt step per outer iteration. This inexact update is used because fully accurate inner solves were found to be unnecessary and sometimes detrimental to practical

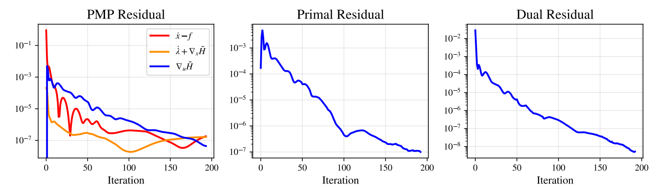

Fig. 2. Outer-iteration residual histories of the inexact PI-ADMM solver on the nominal 3-DoF PDG benchmark. The plotted quantities are the residual of the augmented PMP least-squares problem (22), the ADMM primal residual $\|g(t, x^k, u^k) - z^k\|$, and the ADMM dual residual $\|z^k - z^{k-1}\|$.

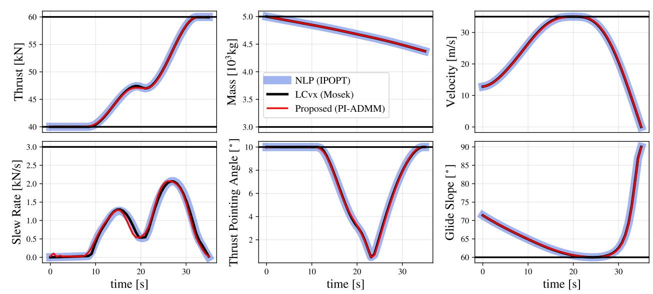

Fig. 3. Constraint-relevant state and control histories for the nominal 3-DoF PDG benchmark obtained by IPOPT, LCvx, and PI-ADMM.

outer-loop progress, consistent with inexact operator-splitting viewpoints [21]. The residuals show a short transient and mild non-monotonicity, but the oscillations remain limited and all three residual measures decrease by several orders of magnitude, indicating improved primal-dual consistency over the iterations.

Fig. 3 compares the state and control profiles obtained by the proposed method and the baseline solvers. All three methods produce closely matching trajectories, and the proposed PI-ADMM solution satisfies the prescribed constraints throughout the horizon. The close agreement in the trajectory profiles indicates that the projection-based constraint handling remains effective even in the presence of the nonconvex thrust-magnitude constraint. This qualitative agreement is also reflected in the quantitative metrics reported in Table II. The three methods yield nearly identical fuel consumption and similarly small constraint violations, all on the order of 10^{-6}. The propagation errors are also comparable, with the proposed PI-ADMM method achieving the smallest value

TABLE II

QUANTITATIVE PERFORMANCE COMPARISON ON THE BENCHMARK.

Method	**Propagation Error**	**Fuel Consumption**	**Constraint Violation**
NLP (IPOPT)	0.109	626.52	1.18×10^{-6}
LCvx (MOSEK)	0.109	626.53	1.72×10^{-6}
Proposed (PI-ADMM)	0.068	626.46	2.30×10^{-6}

*Propagation error denotes the state mismatch obtained by roll-out of the returned solution, fuel consumption is computed as $m_0 - m(t_f)$, and constraint violation denotes the maximum absolute path-constraint residual over the horizon.

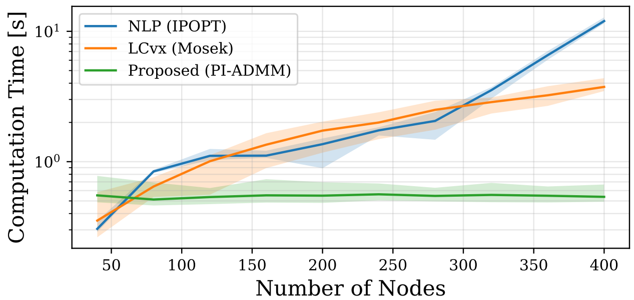

Fig. 4. Wall-clock solve-time scaling with temporal discretization size for the nominal 3-DoF PDG benchmark. For each node count, 50 independent runs are performed for IPOPT, LCvx, and PI-ADMM; solid lines indicate the empirical mean and shaded bands indicate the corresponding variability.

among the tested solvers. These results indicate that the proposed method attains solution quality comparable to both baseline methods while preserving accuracy.

When the solution quality is comparable, the main distinction becomes computational efficiency. Fig. 4 shows the solve-time statistics as the number of temporal nodes increases from 40 to 400 over 50 runs per setting. The proposed PI-ADMM solver is consistently the fastest across almost the entire range of node counts, and its solve time remains nearly flat as the node count increases.

This behavior follows from the coefficient-space formulation of the primal update. Since the Jacobian in (23) satisfies $J \in \mathbb{R}^{Nd_r \times p}$, increasing the number of collocation nodes only enlarges the residual system, while the LM linear system $J^\top J + \epsilon I \in \mathbb{R}^{p \times p}$ remains fixed in size. Consequently, node refinement mainly affects residual evaluation, resulting in a much weaker scaling with the node count than in direct transcription methods.

V. Conclusion

This paper presented PI-ADMM, a physics-informed ADMM framework for constrained continuous-time optimal control, and applied it to a 3-DoF PDG problem. By combining a continuous-time indirect primal update with projection-based ADMM constraint handling, the method preserves PMP structure while effectively enforcing path feasibility. Numerical results showed solution quality comparable to IPOPT and LCvx in terms of fuel consumption, constraint satisfaction, and propagation error, while achieving the shortest solve time among the tested methods.

The main strength of the proposed framework is its ability to combine indirect optimality structure with practical constraint handling in a computationally efficient manner, particularly through a coefficient-space primal update whose dominant linear solve remains fixed in dimension under node refinement. These properties suggest that the method is promising for onboard guidance applications such as planetary landing and reusable launch vehicle recovery. Its current limitations include reliance on empirical parameter tuning and the fact that the present optimality interpretation is most naturally justified under the convex-set assumption, while a general guarantee for nonconvex expansive projections remains unavailable. Future work includes extension to 6-DoF dynamics, free-final-time problems, and stronger theoretical analysis of nonconvex projection-based updates.

References

[1] D. P. Bertsekas, "Nonlinear programming," *J. Oper. Res. Soc.*, vol. 48, no. 3, pp. 334–334, 1997.
[2] L. Biegler and V. Zavala, "Large-scale nonlinear programming using IPOPT: An integrating framework for enterprise-wide dynamic optimization," *Comput. Chem. Eng.*, vol. 33, no. 3, pp. 575–582, 2009.
[3] P. E. Gill, W. Murray, and M. A. Saunders, "SNOPT: An SQP algorithm for large-scale constrained optimization," *SIAM Rev.*, vol. 47, no. 1, pp. 99–131, 2005.
[4] Y. Mao, M. Szmuk, and B. Açıkmeşe, "Successive convexification of non-convex optimal control problems and its convergence properties," in *Proc. IEEE Conf. Decis. Control (CDC)*, 2016, pp. 3636–3641.
[5] B. Açıkmeşe, J. M. Carson, and L. Blackmore, "Lossless convexification of nonconvex control bound and pointing constraints of the soft landing optimal control problem," *IEEE Trans. Control Syst. Technol.*, vol. 21, no. 6, pp. 2104–2113, 2013.
[6] L. Pontryagin, *Mathematical Theory of Optimal Processes*, ser. Classics of Soviet Mathematics. Taylor & Francis, 1987.
[7] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *J. Comput. Phys.*, vol. 378, pp. 686–707, 2019.
[8] D. Mortari, "The theory of connections: Connecting points," *Mathematics*, vol. 5, no. 4, 2017.
[9] R. Furfaro and D. Mortari, "Least-squares solution of a class of optimal space guidance problems via theory of connections," *Acta Astronaut.*, vol. 168, pp. 92–103, 2020.
[10] G.-B. Huang, Q.-Y. Zhu, and C.-K. Siew, "Extreme learning machine: theory and applications," *Neurocomputing*, vol. 70, no. 1-3, pp. 489–501, 2006.
[11] E. Schiassi, A. D'Ambrosio, H. Johnston, M. De Florio, K. Drozd, R. Furfaro, F. Curti, and D. Mortari, "Physics-informed extreme theory of functional connections applied to optimal orbit transfer," in *Proc. AAS/AIAA Astrodynamics Specialist Conf.*, 2020, pp. 9–13.
[12] R. Furfaro, A. D'Ambrosio, E. Schiassi, and A. Scorsoglio, "Physics-informed neural networks for closed-loop guidance and control in aerospace systems," in *Proc. AIAA SCITECH 2022 Forum*, 2022.
[13] A. D'Ambrosio, E. Schiassi, H. Johnston, F. Curti, D. Mortari, and R. Furfaro, "Time-energy optimal landing on planetary bodies via theory of functional connections," *Adv. Space Res.*, vol. 69, no. 12, pp. 4198–4220, 2022.
[14] H. Johnston, E. Schiassi, R. Furfaro, and D. Mortari, "Fuel-efficient powered descent guidance on large planetary bodies via theory of functional connections," *J. Astronaut. Sci.*, vol. 67, no. 4, p. 1521–1552, Sep. 2020.
[15] A. D'Ambrosio, B. Benedikter, and R. Furfaro, "Physics-informed pontryagin neural networks for path-constrained optimal control problems," *J. Guid. Control Dyn.*, vol. 48, no. 8, pp. 1861–1877, 2025.
[16] M. Liu, M. Hou, J. Wang, and Y. Cheng, "Solving two-dimensional linear partial differential equations based on chebyshev neural network with extreme learning machine algorithm," *Eng. Comput.*, vol. 38, no. 2, pp. 874–894, 07 2020.
[17] J. Choi and J.-H. Kim, "Powered descent guidance via first-order optimization with expansive projection," *IEEE Access*, vol. 12, p. 46232–46240, 2024.
[18] G. Park, J. Choi, D. H. Jeong, and J.-H. Kim, "Optimal impact angle guidance via first-order optimization under nonconvex constraints," in *Proc. 2024 Amer. Control Conf. (ACC)*. IEEE, 2024, pp. 778–784.
[19] J. Choi and J.-H. Kim, "LiftProj: Physics-informed Koopman lifting and projection for nonlinear optimal control via first-order optimization," *IEEE Control Syst. Lett.*, vol. 9, pp. 817–822, 2025, also presented at the 64th IEEE Conf. Decis. Control (CDC), 2025.
[20] J. Choi, D. Lee, and J.-H. Kim, "Neural projection operators for real-time 6-DoF powered descent guidance," in *Proc. AIAA SCITECH 2026 Forum*, 2026, p. 1169.
[21] J. Eckstein and D. P. Bertsekas, "On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators," *Math. Program.*, vol. 55, no. 1–3, pp. 293–318, 1992.