Applications of Quantum Optimization Algorithms in Low-Thrust Spaceflight Navigation

Aneesh Sattiraju* and Skyler Kim.†
Georgia Institute of Technology, Atlanta, Georgia, 30332

Low-thrust spaceflight, categorized by low continuous thrust over time, has become a keystone of modern astrodynamics, driving the next generation of space exploration through autonomy, navigability, and efficiency. However, optimizing the trajectory of low-thrust spacecraft presents significant challenges; their trajectories evolve in an environment dominated by complex gravitational interactions and chaotic nonlinear dynamics where small errors compound over time, affecting mission success. Current classical trajectory design methods excel at refining short-range trajectories through local search, but struggle in understanding and traversing the global landscape of multi-body dynamics. Consequently, classical methods often yield trajectories that seem plausible but are infeasible over time. This paper advances the splitting of spaceflight navigation into search and refinement–integrating quantum optimization and machine learning alongside classical solvers– aiming to address the lack of investigation into global search methods. By investigating the unique applications of the Quantum Approximate Optimization Algorithm and Residual Networks for global pathway discovery, and classical solvers for refining each path into executable trajectories, this conceptual hybrid architecture aims to catalyze improvements in robust trajectory design for chaotic environments.

I. Introduction and Motivation

Modern spacecraft trajectory design requires maintaining precise motion through chaotic gravitational environments that resist simple optimization. Unlike powerful, impulsive burns that can be modeled as discrete "kicks" to velocity [1], a low-thrust spacecraft gradually shapes velocity, enabling missions requiring trajectory control over long durations of time, such as station keeping, orbit maintenance, and orbital shaping. Low-thrust propulsion systems are favorable compared to chemical propulsion systems in long-horizon missions due to high fuel efficiency and high specific impulse (I_{sp}) [2]. This capability is expected to play a vital role in future deep space missions, including Mars transfers and asteroid belt navigation. Since low-thrust navigation relies on sustained guidance rather than isolated maneuvers, trajectory accuracy becomes highly sensitive to any small deviations that may amplify over time due to surrounding nonlinear dynamics [3]. As a result, trajectory optimization extends beyond an optimal control problem; it requires searching for a structured, stable pathway within a complex environment. Conventional gradient-based low-thrust optimization techniques are designed to refine trajectories once an initial path is given; however, these classical methods are highly dependent on the quality of the initial guess and often converge to local, suboptimal solutions [4].

The actual discovery of globally viable pathways for trajectory optimization remains computationally challenging. Established global metaheuristic optimization algorithms, such as differential evolution, genetic algorithms, and particle swarm algorithms, have demonstrated strong capability in exploring complex optimization landscapes and avoiding local minima. However, these methods require full, explicit cost evaluation of candidate solutions to guide search [5]. Low-thrust trajectories are discretized into numerous time nodes with a thrust vector optimized at each node, naturally rendering low-thrust optimization as a high-dimensional control problem, which is, principally, dimensionally infinite. As the number of time nodes increases, low-thrust control problems obtain an immense amount of decision variables. In general, in high-dimensional search spaces, the number of required evaluations for complete cost estimation would grow substantially, generating a significant computational burden for metaheuristic algorithms [6]. This curse of dimensionality constrains even robust metaheuristic algorithms where parameters exceed thousands of decision variables, limiting convergence speed and precision [6].

Given the extreme sensitivity of navigation in multi-body dynamics, reliance on locally guided search and computationally expensive metaheuristic approaches alone could lead to fragile convergence behavior, prompting investigation into alternative search structures capable of efficient cost evaluation. Motivated by these structural

*Undergraduate Student Member, Georgia Institute of Technology, AIAA Student Member 1861038
†Undergraduate Student Member, Georgia Institute of Technology, AIAA Student Member 1861503

challenges, this work separates trajectory search and refinement, exploring quantum optimization approaches as a global search method within a proposed hybrid framework for low-thrust trajectory design. In particular, the Quantum Approximate Optimization Algorithm (QAOA) is examined as a combinatorial search procedure that may integrate with classical refinement methods. While the proposed hybrid framework is conceptual rather than mission-validated, this perspective aims to characterize how alternative optimization techniques may broaden trajectory design methods to operate effectively in chaotic environments.

The remainder of this paper is organized as follows: Section II reviews background and foundational concepts from low-thrust spaceflight and optimization approaches; Section III introduces the proposed surrogate-assisted hybrid quantum-classical architecture; Section IV discusses potential applications and metrics; and Section V concludes with a discussion and future directions.

II. Background and Foundation

A. Low-Thrust Spaceflight

The surrounding environment of a low-thrust spacecraft is complex; the spacecraft's thrust acceleration is relatively weak compared to the dominant gravitational forces in the surrounding multi-body region. Consequently, low-thrust trajectory optimization naturally relies on control-based structures such as state-vectors and state-space. The spacecraft's state at any given time t evolves within a seven-dimensional state space, and is traditionally modeled by its state vector:

$$x(t) = [\,\mathrm{r}(t), \mathrm{v}(t), m(t)\,] \in \mathbb{R}^7 \tag{1}$$

with $m(t)$ representing mass, and $r(t)$ and $v(t)$ representing the three-dimensional position and velocity of the spacecraft, respectively [7]. Thrust acts as the input that controls spacecraft motion, and is modeled by the control variable, the thrust vector:

$$\mathrm{u}(t) = T(t)\,\hat{\mathrm{u}}(t) \tag{2}$$

with $T(t)$ representing thrust magnitude, and $\hat{u}(t)$ representing direction. Together, the state vector defines the spacecraft's current condition, and the control variable determines how the system will evolve over time. The derivative of the state vector $x(t)$ gives the system dynamics:

$$\dot{x}(t) = f\big(x(t), u(t), t\big) = \begin{bmatrix} \mathrm{v}(t) \\ \mathrm{a}(t) \\ \dot{m}(t) \end{bmatrix} \tag{3}$$

where $\dot{m}(t)$ captures propellant depletion, and $a(t)$ represents the combined influence of surrounding gravitational forces and thrust input. This state-space form depicts the spacecraft's condition and evolution given an initial state $x(t)$ and thrust profile $u(t)$, making it a core component of low-thrust optimization [8].

B. Dynamical Systems Theory in Low-Thrust Spaceflight

Deep space missions rarely operate under the gravitational influence of a single body; low-thrust missions must navigate through complex gravitational fields and multi-body dynamics [9], making prediction and optimization substantially chaotic and harder than classical two-body physics. Therefore, recognizing how systems evolve over time through an aerospace-oriented Dynamical Systems Theory (DST) perspective is essential for understanding the constraints and demands of trajectory optimization.

As discussed in the previous section, a spacecraft evolves in a state space defined by its position, velocity, and mass. The nonlinear structure of this space reveals gravitational features that are not apparent in physical space yet drastically influence motion. When a spacecraft's state is examined alongside the geometric structures it operates within, its state space is often referred to as phase space, revealing the underlying structures that determine motion. Plotting a state vector over time in a spacecraft's phase space yields a precise path of motion, known as the trajectory. Because spacecraft operate in nonlinear environments, the geometry of phase space heavily constrains feasible low-thrust trajectories, and any optimization method must engage deeply with state-vectors and phase space trajectories [10].

1. *Fixed Points and Equilibria*

Fixed points or equilibrium points are where the system exhibits no change; the next state is the same as the current state (state derivative $\dot{x}$(t) is 0). At these points, net forces and acceleration cancel relative to the motion of the spacecraft, resulting in a constant state vector. Fixed points determine the stability of a trajectory through stable and unstable equilibria. Locations where a system naturally returns to its original state and attracts nearby trajectories are known as stable equilibria, whereas points where small deviations exponentially amplify and repel nearby trajectories are known as unstable equilibria. Gravitational models and coordinate transformations reveal regions with different equilibria, and trajectories will naturally approach stable equilibria and diverge away from unstable equilibria, creating regions of sensitive dynamical behavior [11].

2. *Lagrange Points and Manifolds*

In the Circular Restricted Three-Body Problem (CR3BP), five critical equilibrium solutions, known as Lagrange points (L1-L5), emerge in the rotating frame of the two primary bodies orbiting their common center of mass. Net acceleration of the third negligible-mass body, potentially a spacecraft, vanishes relative to the frame of reference, allowing a spacecraft to remain stationary at these points relative to the primary bodies [12]. L4 and L5 are stable equilibria with restoring dynamics that keep spacecraft in these regions, while L1-L3 are unstable equilibria in which tiny deviations grow exponentially. Spacecraft near unstable equilibria, such as the L1-L3 regions, require active station-keeping thrust to maintain their position, forcing spacecraft to instead utilize halo orbits, or small looping paths that reduce control effort while remaining in the dynamical structure. At stable L4-L5 regions, constant thrust is not necessary. The coexistence of attracting and repelling equilibria consequently organizes motion into invariant manifolds - geometric structures or pathways derived from the underlying gravitational environment that determine the transport through phase space [9]. The exploitation of these structures enables spacecraft to leverage gravitational geometry rather than sustained thrust. A spacecraft may enter a halo orbit, slip off an unstable manifold, get captured by a stable manifold, and experience transportation between regions with minimal fuel cost. CR3BP models are applied to locally dominant two-body gravitational systems (e.g., Earth-Moon, Sun-Earth, Sun-Jupiter), each of which admits their own Lagrange points under this approximation. In this sense, long-duration deep space missions may be interpreted as a sequence of transitions between locally dominant CR3BP subsystems, and the associated invariant manifolds provide structured transport mechanisms within the broad multi-body environment. Exploiting dynamical architecture through local manifold-guided motion is a foundational approach to deep space navigation and trajectory optimization, constraining the set of feasible trajectories.

3. *Chaos, Perturbations, and Sensitivity to Initial Conditions*

In dynamical systems, chaos is deterministic behavior with extreme sensitivity to initial conditions. Physics ensures no randomness is involved, but due to complex multi-body dynamics, two nearly identical initial states may diverge into completely different futures [13], explaining how centimeter-scale deviations eventually produce trajectory errors spanning thousands of miles. Additionally, gravitational dynamics are inherently nonlinear. In a two-body system, gravitational motion is integrable and predictable; however, the introduction of additional bodies renders the equations non-integrable, producing coupling between gravitational influences that prevents the analytical isolation of variables. As a result, small deviations accumulate, amplifying trajectory sensitivity and limiting long-horizon predictability.

Perturbations introduce small deviations to an ideal system, influencing its equilibrium state and causing nearby trajectories to diverge over time. These perturbations emerge not only from numerous gravitational forces, but also from solar radiation pressure, thrust execution errors, and imperfect state estimations since exact position and velocity are never known. The lack of smooth behavior and the high chaotic nature of a spacecraft's phase space reveal the fundamental limitations in current optimization methods when applied to low-thrust spaceflight. Effective trajectory design must account for chaos, exploit gravitational geometry, and reason over complex dynamical structures to discover stable long-term trajectories.

C. Classical Optimization Methods and Limitations

Low-thrust trajectory optimization is typically treated as a continuous optimal control problem and frequently addressed using gradient-based methods [14]. The setup for low-thrust spaceflight is different from impulsive calculations, because the optimal decision is a long-horizon steering law rather than a single burst. Each optimization function contains four ingredients:

1) State Vector
2) Control variable
3) Dynamics
4) Cost

Typical cost function examples aim to address fuel usage, flight time, or weighted trade-offs (fuel-time, fuel-risk, etc.). Trajectory optimization seeks to minimize this cost [15]. However, classical methods are designed to refine an initial trajectory guess that is given rather than discover one from scratch. In practice, candidate trajectories are constructed through analytical approximations, continuation strategies, physics-informed heuristics, or past mission trajectories prior to refinement [4, 16].

1. Classical Methods

Optimization problems are often approached either indirectly or directly. The classic indirect approach, known as Pontryagin's Minimum Principle (PMP), is a foundational control theory method that demonstrates the conditions an optimal thrust profile must satisfy without actually computing the trajectory itself. This is done by encoding the sensitivity of how much the cost is affected by changing the state into a variable called the costate [17]. The goal of PMP is to determine the thrust that leads to the best overall outcome at the end of the mission, assuming the future behaves smoothly, prioritizing long-horizon optimality over immediate favorability. Consequently, PMP enables highly precise trajectory refinement once a viable path has been confirmed. The primary direct approach, known as Nonlinear Programming (NLP), operates by initializing a complete candidate trajectory and iteratively refining it using gradient-based optimization until convergence at an optimal solution [14]. Transcribing the continuous control problem over the entire time-horizon by discretizing the state and control parameters over several time intervals enables the entire trajectory to be optimized simultaneously. Importantly, both PMP and NLP operate directly on continuous dynamics, enforcing the equations of motion as constraints and reducing convergence at physically infeasible solutions. Both indirect and direct approaches, therefore, provide precise trajectory refinement once a viable candidate is available [8].

2. Structural Limitations

Classical optimization methods rely on local search, primarily leveraging local information to make decisions, and therefore cannot perceive the geometry of the entire cost landscape. Fundamentally, classical optimization unfolds within the basin of attraction surrounding the initial guess, and refinement success highly depends on the quality of the initialized trajectory and the geometry of the cost landscape. The global minimum of the cost function corresponds to the optimal solution for the control problem. Thus, classical solvers operate better with smoother cost landscapes, as demonstrated by local PMP and NLP optimization; optimizers operate by following descending directions, and the steepest descending directions within the landscape are determined using gradients and derivatives, necessitating continuity and smoothness [14].

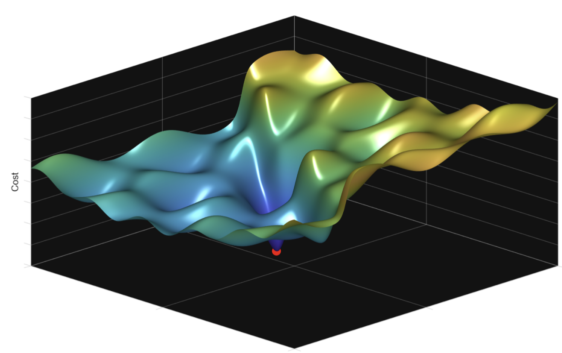

Fig. 1 Illustrative example of a simple nonconvex cost landscape with many local minima and a global basin.

However, the cost landscape of low-thrust spaceflight is nonconvex, containing multiple local extrema which amplify the difficulty of identifying the true minimum. Chaotic regions of space cause tiny changes in thrust to lead to huge changes in trajectory evolution, and the gravitational influences of equilibria render the cost function unpredictable.

The inherent limitation of classical solvers then emerges when confronted with a discontinuous and complex variable landscape; perturbations and chaos resist gradient-based analysis, classical solvers often converge to suboptimal minima, and these algorithms struggle when attempting to cross the chaotic separatrices where the trajectory through phase space becomes increasingly discontinuous. The central challenge in low-thrust trajectory optimization is therefore not trajectory refinement, but the identification of globally viable trajectories in the first place.

D. Quantum Optimization and the QAOA

Overcoming the chaos that impedes effective trajectory design may require global search methods, capable of evaluating a multitude of trajectory alternatives at once rather than the refinement of a single volatile option. Unlike classical algorithms that rely on gradients and individual iterations, quantum optimization offers the potential to explore nonconvex cost landscapes more globally by exploiting the flexibility of qubits to encode a superposition of many candidate solutions in parallel [18, 19].

Quantum-based global search is facilitated by first encoding all variables into one optimization landscape. When the cost function is applied, the selection of which parameter to optimize shapes the entire landscape into a specific configuration, with optimal combinations existing at the lowest energy points. The global potency of quantum methods arises not only from a representation of many candidate solutions simultaneously, but also because they are not confined to local derivative information to decide where to guide search. Accordingly, this paper shifts focus to a programmable hybrid quantum-classical global search framework adept in combinatorial decision making: the Quantum Approximate Optimization Algorithm (QAOA).

The QAOA is a hybrid quantum-classical, or variational circuit, designed originally to find near-optimal solutions for complex combinatorial optimization problems [19] such as portfolio optimization, finance, or logistics, where classical methods struggle due to rugged optimization landscapes. The global search behavior of this algorithm relies on two Hamiltonian functions, the cost Hamiltonian H_{cost} and the mixing Hamiltonian H_{mix} [19], shown in a simplified single-layer QAOA ansatz:

$$|\psi(\gamma,\beta)\rangle = e^{-i\gamma H_{\text{cost}}} e^{-i\beta H_{\text{mix}}} |+\rangle^{\otimes n} \tag{4}$$

where $|+\rangle^{\otimes n}$ ensures a quantum state with an equal superposition over all qubits. Together, these Hamiltonians encode the problem's cost structure and induce exploration across the entire configuration space.

1. Quantum Process

The quantum process begins with all candidate configurations represented simultaneously within a uniform superposition. The cost function is then applied, which indicates the parameter that must be optimized. This encoded information, facilitated by H_{cost}, known as phase, embeds each configuration with a phase proportional to its own cost, encoding the unique geometric nonconvexity of the cost function onto the landscape. At this stage, only information is assigned; while no solution has been determined, H_{cost} embeds the cost function into the entire landscape, which is necessary for subsequent optimization. A mixing layer using H_{mix} is applied, causing the phases to interfere and result in different amplitudes or results of that quantum interference. Constructive interference seeks to amplify low-cost candidates, while destructive interference suppresses high-cost ones, redistributing probability mass to concentrate in low-cost regions [19]. Then, the quantum computer measures the resultant quantum state and outputs a set of candidate configurations (bitstrings) into a classical computer. As probability has been redistributed across optimal low-cost regions, the probability that a selected bitstring was output is proportional to its influence. By embedding cost into phases, the optimization problem is transformed from a topographical search to a wave interference problem, reshaping the search process towards globally optimal regions. After many iterations, the goal of QAOA is to filter out suboptimal minima, sampling bitstrings that may survive local traps and are more likely to be globally optimal.

2. Classical Process

The programmability of QAOA arises from the role of the classical computer: optimizing the parameters that determine the application of H_{cost} and H_{mix} on the quantum circuit. The strength of their interaction determines how the quantum state explores the solution space [20]. The classical computer, using H_{cost}, evaluates the expected cost produced by the quantum state by sampling bitstrings and assigning scalar scores based on their cost. Options with lower scores (lower cost) are favored, and the performance evaluations are sent back to the quantum computer to adjust the parameters (γ, β). By tuning the variational parameters γ and β, the algorithm influences the interaction between H_{cost} and H_{mix} to guide quantum interference to lower-cost regions of the solution space. This process, known as the variational loop, runs repeatedly, evolving the QAOA to increase the probability of sampling high-quality solutions consistent with the global optimum. In the complete QAOA circuit, the quantum computer guides classical computational effort to regions demonstrating global promise and lower cost, enabling global exploration of complex optimization landscapes. The following section examines how these global capabilities may be integrated with classical refinement to promote robust low-thrust trajectory design.

III. Hybrid Quantum-Classical Architecture for Low-Thrust Trajectory Optimization

Although low-thrust motion is governed by continuous differential equations, trajectory design under multi-body environments is inherently combinatorial; the phase-space contains discrete dynamical structures revealed under dominant gravitational systems: fixed points, Lagrange points in different CR3BP systems, invariant manifolds, and capture basins. A spacecraft does not merely flow smoothly through space; efficient navigation demands a choice of how to interact with these regimes - when to enter, leave, exploit, or drift along these structures. From a mission-design perspective, a complete low-thrust mission therefore involves a sequence of discrete dynamical decisions, even if the physical transition between each region is continuous. Classical gradient-based solvers excel in refining this continuous transition once a globally viable pathway has been identified; however, the success of local methods depends on the quality of their initialization, rendering global search as a prerequisite to effective refinement. Furthermore, low thrust sequencing is a permutation problem [21]; for a mission involving n candidate targets or regions, the number of possible visitation sequences scales factorially ($n!$), exploding the decision space for even modest missions. In other words, the greatest challenge of low-thrust trajectory optimization is not computing the physics through a region, but the combinatorial decision of which regions to traverse. As a result, low-thrust trajectory design exhibits a dual structure: a combinatorial search layer for global pathway selection, and a continuous refinement layer for physically executing the trajectory, leading to the hybrid quantum-classical framework developed by this work to address both layers simultaneously.

A. QAOA for global pathway selection

Since the QAOA was designed for finding approximate global solutions for hard combinatorial problems with nonconvex cost landscapes and coupled variables, QAOA naturally emerges as a candidate for the optimization of pathway selection for low-thrust spaceflight. A distinguishing feature of the QAOA is that n-qubit quantum state resides in a Hilbert space of dimension 2^n. Because each qubit represents a two-level system ($|0\rangle$ or $|1\rangle$), there are 2^n computational basis states. Accordingly, the quantum state may represent a superposition over all 2^n decision configurations simultaneously. In contrast to global metaheuristic algorithms that scale with population size and require explicit cost evaluation of each candidate, QAOA offers a fundamentally different search paradigm: redistributing probability over the entire high-dimensional decision space via phase encoding and mixing operations.

Within the proposed framework, each decision regarding interaction with phase-space structures may be encoded onto qubits as binary variables. For example, 1 may correspond to choosing an option (enter resonance orbit) and 0 may correspond to ignoring an option (do not enter). Multiple qubits can represent multiple-choice decisions, such as:

$$
\begin{aligned}
|00\rangle &\rightarrow \text{remain in current dynamical region} \\
|10\rangle &\rightarrow \text{evade a gravitational basin} \\
|01\rangle &\rightarrow \text{approach local invariant manifold (e.g., Sun–Earth } L_4 \text{ Manifold)} \\
|11\rangle &\rightarrow \text{exploit stable transport pathway}
\end{aligned}
$$

The result is a sample bitstring that represents a skeletal trajectory through space:

$$|1\,0\,1\,1\,0\,1\ldots\rangle \tag{5}$$

At this stage, the quantum state is $|+\rangle^{\otimes n}$, meaning every skeleton has an equal amplitude. Then H_{cost} aggregates multiple mission objectives such as time, risk, and propellant into one weighted sum, embedding the cost structure of the optimization landscape into the quantum state:

$$H_{\text{cost}} = \omega_1 H_{\text{fuel}} + \omega_2 H_{\text{risk}} + \omega_3 H_{\text{time}} + \ldots \tag{6}$$

with ω representing the weight, or relative importance of each consideration. As the cost function is applied, the quantum state is shaped, and each sequence now has a phase that is proportional to its cost under the mission objectives. Finally, H_{mix} transforms the phase differences into interference that redistributes the probability mass to lower energy regions, increasing the probability of outputting a globally optimal bitstring. Measurement of the quantum state at the end of the circuit yields globally informed, candidate trajectory skeletons, transforming pathway selection into a structured global search process.

Though QAOA operates on discretized commitments, the physics of low-thrust motion is continuous. The algorithm is not intended to represent the physics that make the specific trajectory output possible; instead, its role is to output a globally favorable trajectory skeleton by operating on distinct decisions. In chaotic environments, global pathway decisions largely determine mission success and the success of subsequent local refinement. Therefore, the proposed architecture relies on quantum optimization to search the combinatorial navigation of low-thrust spaceflight and then use the proposed skeleton as an initialization for established classical methods to refine into an executable trajectory.

B. Classical Solvers in the Hybrid Architecture

1. Cost Evaluation in the Variational Loop

Within the QAOA loop, the quantum algorithm will output candidate bitstring plans that are evaluated holistically as complete mission configurations rather than a sequence of decisions. A classical optimizer computes the expected cost of the quantum state and updates the QAOA's parameters (γ, β), looping this process to iteratively shape the probability distribution towards global regions. However, in standard QAOA, the scoring of each option assumes that the objective cost is static and can be computed efficiently and inexpensively for each sampled bitstring [20]. Low-thrust spaceflight optimization struggles with this assumption, as accurately computing mission cost requires modeling nonlinear dynamics, long-horizon sensitivity analysis, and stability assessments [3, 8]. The primary bottleneck is therefore not quantum sampling, but the computational expense of computing the expected cost in the variational loop. Addressing this bottleneck requires a fast, trainable surrogate capable of learning and representing nonlinear dynamics for cost evaluation [22].

2. Trajectory Refinement

The conclusion of the variational loop results in a small set of globally promising trajectory skeletons, completing the role of QAOA. Classical optimal control methods then assume responsibility for transforming trajectory skeletons into executable trajectories. Each skeleton becomes a mathematical constraint [14, 17], encoded as trajectory constraints in NLP or boundary conditions for PMP. In direct approaches, the proposed trajectory is discretized over time, and the thrust profile $T(t)$ and state evolution $x(t)$ are solved under the governing dynamics. Indirect approaches derive the necessary conditions for the proposed trajectory through costate variables. In both methods, the classical solver complements the QAOA by determining the continuous control required to realize the discrete mission plan. Essentially, QAOA finds where to go, while classical control methods find out how.

Although classical refinement relies on local, gradient-based methods, the dominant failure mode in low-thrust navigation is choosing the wrong global pathway, not refining it [14]. Instead of refining locally through the full phase space, classical solvers are meaningfully initialized and local optimization is confined to globally informed regions when complemented with QAOA.

C. ResNet Model as the Enabling Component

Residual Networks (ResNets) are introduced in this model as trainable surrogates capable of rapidly estimating trajectory cost and structure. They are warranted in this context due to their stability in deep, nonlinear systems, and because their architecture is designed to learn corrections from an existing baseline rather than absolute mappings [23]. This strategy aligns with the refinement-based nature of trajectory optimization, alleviating the degradation problem observed in deep plain networks [23]. The use of skip connections further alleviates vanishing gradient effects, enabling stable training in nonlinear physical spaces where long-horizon behavior is difficult to model.

1. ResNet as a Surrogate Model for Cost Evaluation

Accurate parameter updates within the QAOA loop depend on reliable evaluation of the sampled trajectory cost. However, in low-thrust environments, computing the expected cost depends on highly sensitive, chaotic dynamics, and direct computation is often too computationally expensive to perform within the variational loop. Addressing this limitation may require a cost model capable of learning the dynamical environment. Accordingly, catalyzing the QAOA loop with ResNets enables rapid approximate trajectory evaluations [22]. These networks, trained on representative trajectory data and physics in non-Euclidean phase space, reduce the computational burden of cost estimation while remaining faithful to the governing dynamics. Operationally, this is done through a system called a "Surrogate-Assisted Variational Loop" [24], where the QAOA generates a quantum state and the ResNet provides near-instant cost estimates, facilitating efficient parameter updates without full physical computation. In this model, ResNets that act as surrogate cost models may also filter out globally informed but structurally poor, risky, and infeasible trajectories, increasing solution quality.

2. Initialization of Classical Control Methods

ResNets may also support the stability of classical refinement for trajectory optimization. Even when QAOA has identified a globally favorable trajectory skeleton, indirect methods are highly sensitive to initialization and direct methods can converge to physically valid but topologically inconsistent trajectories [8]. ResNets would not replace NLP or PMP methods, but provide educated initial conditions [16, 22] that reduce numerical sensitivity, add marginal corrections, and promote convergence within the dynamical basin defined by the trajectory skeleton. Trained on families of trajectories and phase space structures, the ResNets can initialize solvers with approximate thrust histories, durations, and timings, as well as approximate costate estimates that improve PMP convergence - a primary issue revolving around indirect methods. By guiding classical solvers into safer regions, gradients and refinement retain their meaning, and the chances of converging to incorrect minima are reduced.

D. Quantum-Classical Surrogate Variational Loop

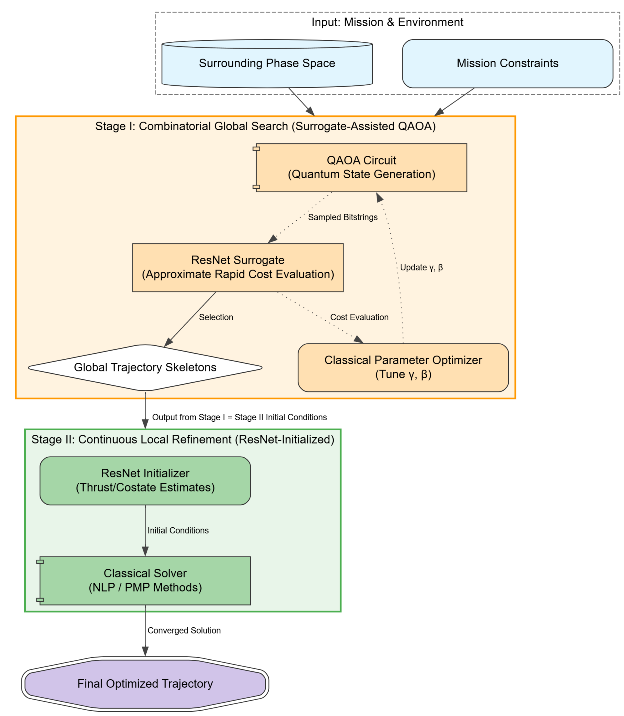

Fig. 2 Proposed Hybrid Quantum-Classical Architecture.

The proposed framework for low-thrust spaceflight trajectory optimization relies on three main components. Each component addresses a distinctive failure mode of low-thrust navigation by treating low-thrust optimization as continuous in physics but combinatorial in structure.

1) QAOA addresses the difficulty of global pathway selection by searching for a globally coherent sequence of commitments

2) Trained and phase space informed ResNets enable efficient cost evaluation within the variational loop to provide stable parameter updates until the loop converges. ResNets also support initialization for classical control methods for the final output global bitstrings
3) Classical solvers perform refinement by solving the continuous control problem under the governing dynamics to obtain a complete trajectory

Together, these three elements aim to form a hybrid architecture capable of integrating global search with physically consistent refinement to enable more robust trajectory design.

IV. Potential Mission Applications and Validation Methods

Although the proposed framework is conceptual and offers an alternative model for global navigation, the transition from a theoretical architecture to operational implementation depends on clear mission applications and a defined validation strategy.

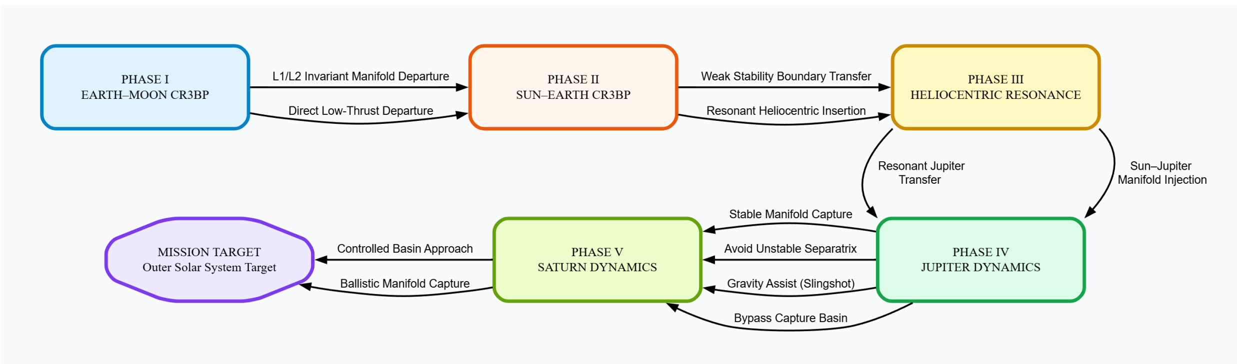

Fig. 3 Example of discrete trajectory commitments in a long-horizon outer Solar System transfer.

A. Representative Application

A representative mission application for this architecture may involve a long-horizon transfer from Earth to the outer Solar System. As illustrated in Fig. 3, the mission demands sequential interaction with multiple gravitational regions in which well-established, valid transport mechanisms are available. These mechanisms are part of the larger Interplanetary Transport Network, a framework that enables low-energy transport by exploiting manifolds and related dynamical structures [25]. However, the same unstable manifolds that enable low-energy transport also introduce sensitive, chaotic behavior near dynamical boundaries. Therefore, the primary challenge arises not from individual maneuvers, but from selecting a consistent, efficient set of region-to-region commitments. The proposed architecture may be particularly well-suited for the context described above because it addresses the structural difficulty of trajectory sequencing. Each transition between gravitational regions, whether via manifolds, gravity assists, or a separatrix evasion, is a discrete decision by nature. As the number of target regions and transport mechanisms increase, the amount of possible transition sequences grows exponentially, challenging local search methods that refine within the basin of a single initialized trajectory. Furthermore, population-based metaheuristic algorithms must explicitly evaluate candidate trajectories with high computational cost; however, the proposed architecture operates by separating selection from refinement, enabling combinatorial reasoning over families of trajectory skeletons before demanding CPU-intensive numerical integration. In long-horizon missions, the cost of evaluating each trajectory option compounds, motivating the integration of surrogate-assisted cost evaluation for rapid quantum parameter updates to reduce the computational burden of the proposed framework. Finally, because classical optimal control methods are applied only to globally informed candidates, convergence to suboptimal basins may become less likely, potentially allowing the overall trajectory optimization framework to scale with mission complexity while maintaining computational efficiency.

B. Validation Criteria

Validation of the proposed framework would require a structured comparative evaluation against established trajectory design methods across long-horizon multi-body missions. A set of sequences is considered globally favorable if they correspond to distinct region-to-region commitments, yield lower mission costs, and remain dynamically feasible when subject to refinement. In classical evolutionary algorithms, global search (exploration) and local refinement (exploitation) are intertwined within the same evolutionary cycle rather than decoupled into structurally distinct stages [26]. In this framework, global exploration occurs at the combinatorial level prior to continuous optimization, allowing distinct pathways to be compared before commitment. Metrics for benchmarking the performance of low-thrust optimization techniques may include propellant consumption, flight time, computational runtime, and convergence rate. Additionally, search criteria such as sensitivity to state perturbations and convergence under discontinuous cost landscapes would offer additional insight into algorithmic performance. Relating to refinement, assessment of different initialization strategies may involve comparing classical heuristic-based initialization against the proposed ResNet-informed initialization. Metrics such as convergence reliability, number of solver iterations, final cost, and number of failed refinement attempts could clarify if surrogate-informed initialization improves stability and reduces sensitivity to poor initial guesses. In this context, an optimization architecture may be considered advantageous if it demonstrates a higher probability of identifying diverse, low-cost, feasible trajectory skeletons; reduced computational burden; and an increased convergence reliability as mission complexity increases.

The proposed framework was structured with these objectives in mind; however, current term implementation remains constrained since QAOA currently operates within the Noise Intermediate-Scale Quantum (NISQ) era [27]. Measurement noise, gate errors, and decoherence may cause quantum states to drift from their intended evolution, reducing the reliability of sampled bitstrings [28]. Additionally, representing chaotic phase-space decisions may require more qubits and circuit depth than current NISQ quantum hardware supports. As a result, advances in qubit density, error mitigation, and coherence are needed before advanced quantum optimization may be realized. Near-term implementation of the proposed architecture may rely on classical emulation of quantum circuits and hybrid approximations while full-scale quantum hardware continues to develop.

V. Conclusions and Future Directions

This work advances the perspective that low-thrust trajectory optimization is most naturally treated as a dual problem of search and refinement, motivating the integration of quantum optimization in a domain traditionally confined to classical methods. By exploring the applications of QAOA, an algorithm designed for combinatorial optimization, to low-thrust spaceflight, this work proposes a hybrid quantum-classical architecture that separates pathway discovery from execution. Quantum processes evaluate discretized mission commitments, embedded ResNets enable rapid cost estimation within the variational loop and support the initialization of the classical solvers that refine a trajectory skeleton into an executable path. Collectively, this structure forms an adaptive and scalable optimization engine in which quantum methods determine where to go and classical methods determine how.

Future work should focus on operationalizing these components in progressively realistic conditions, including advances in noise and error mitigation, preservation of meaningful quantum interference, and the number of reliable qubits for optimization. Exploration of methods for encoding continuous control variables into a quantum state may also further reduce dependence on classical refinement. As trained surrogate models mature, onboard autonomy may expand to support real-time trajectory adaptation without excessive reliance on ground-based intervention.

By recognizing low-thrust trajectory optimization as both combinatorial in structure and continuous in physics, future navigation systems may no longer merely compute trajectories, but also reason over space itself, harnessing chaos not as a barrier, but as a medium for navigation.

References

[1] Bate, R. R., Mueller, D. D., and White, J. E., *Fundamentals of Astrodynamics*, Dover Publications, New York, 1971.

[2] Parrish, N. L., and Scheeres, D. J., "Low-Thrust Trajectory Optimization with Simplified SQP Algorithm," Tech. Rep. NASA GSFC-E-DAA-TN45591-2 (AAS 17-832), NASA Goddard Space Flight Center, Greenbelt, MD, 2017.

[3] Morante, D., Sanjurjo Rivo, M., and Soler, M., "A Survey on Low-Thrust Trajectory Optimization Approaches," *Aerospace*, 2021. https://doi.org/10.3390/aerospace8030088.

[4] Kelly, M., "An Introduction to Trajectory Optimization: How to Do Your Own Direct Collocation," *SIAM Review*, 2017. https://doi.org/10.1137/16M1062569.

[5] Bozovic, A., Jovanovic, L., Dobrojevic, M., et al., "Exploring the Applicability of Decision Trees and Deep Neural

Networks Optimized by Metaheuristics for Predictive Maintenance in Milling," *Journal of Supercomputing*, 2025. https://doi.org/10.1007/s11227-025-08082-0.

[6] Khosrowshahli, R., Rahnamayan, S., and Ombuki-Berman, B., "Massive Dimensions Reduction and Hybridization with Metaheuristics in Deep Learning," *Proceedings of the 2024 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE)*, IEEE, 2024, pp. 469–475. https://doi.org/10.1109/CCECE59415.2024.10667313.

[7] Longuski, J. M., Guzmán, J. J., and Prussing, J. E., *Optimal Control with Aerospace Applications*, Springer, New York, 2013.

[8] Rao, A. V., "A Survey of Numerical Methods for Optimal Control," *Advances in the Astronautical Sciences*, Vol. 135, American Astronautical Society, San Diego, CA, 2009, pp. 497–528.

[9] Koon, W. S., Lo, M. W., Marsden, J. E., and Ross, S. D., *Dynamical Systems, the Three-Body Problem, and Space Mission Design*, California Institute of Technology, Pasadena, CA, 2011.

[10] Parker, J., and Anderson, R., *Low-Energy Lunar Trajectory Design*, JPL Deep Space Communications and Navigation Series, Vol. Monograph 12, Wiley, Hoboken, NJ, 2013.

[11] Wiggins, S., *Introduction to Applied Nonlinear Dynamical Systems and Chaos*, 2nd ed., Springer, New York, 2003.

[12] Gómez, G., Llibre, J., Martínez, R., and Simó, C., *Dynamics and Mission Design Near Libration Points, Vol. I: Fundamentals*, World Scientific Publishing, Singapore, 2001.

[13] Belbruno, E., *Capture Dynamics and Chaotic Motions in Celestial Mechanics*, Princeton University Press, Princeton, NJ, 2004.

[14] Betts, J. T., *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*, 2nd ed., Society for Industrial and Applied Mathematics, Philadelphia, PA, 2010.

[15] Hull, D. G., *Optimal Control Theory for Applications*, Springer, New York, 2003.

[16] Sinha, A., and Beeson, R., "Initial Guess Generation for Low-Thrust Trajectory Design with Robustness to Missed-Thrust Events," *Journal of Guidance, Control, and Dynamics*, 2025. https://doi.org/10.2514/1.g009050.

[17] Naidu, D. S., "Pontryagin Minimum Principle," *Optimal Control Systems*, CRC Press, Boca Raton, FL, 2003, pp. 249–291.

[18] Nielsen, M. A., and Chuang, I. L., *Quantum Computation and Quantum Information*, 10th ed., Cambridge University Press, Cambridge, UK, 2010.

[19] Farhi, E., Goldstone, J., and Gutmann, S., "A Quantum Approximate Optimization Algorithm," *arXiv*, 2014. https://doi.org/10.48550/arXiv.1411.4028.

[20] Zhou, L., Wang, S., Choi, S., et al., "Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices," *Physical Review X*, 2020. https://doi.org/10.1103/PhysRevX.10.021067.

[21] Vasile, M., Minisci, E., and Locatelli, M., "Analysis of Some Global Optimization Algorithms for Space Trajectory Design," *Journal of Spacecraft and Rockets*, 2010. https://doi.org/10.2514/1.45742.

[22] Di Fiore, F., and Minisci, E., "Surrogate Models for Optimization of Dynamical Systems," *arXiv*, 2021. https://doi.org/10.48550/arXiv.2101.10189.

[23] He, K., Zhang, X., Ren, S., and Sun, J., "Deep Residual Learning for Image Recognition," *arXiv*, 2015. https://doi.org/10.48550/arXiv.1512.03385.

[24] Shaffer, R., Kocia, L., and Sarovar, M., "Surrogate-Based Optimization for Variational Quantum Algorithms," *Physical Review A*, 2023. https://doi.org/10.1103/PhysRevA.107.032415.

[25] Lo, M., "The InterPlanetary Superhighway and the Origins Program," *Proceedings of the IEEE Aerospace Conference*, Vol. 7, IEEE, 2002, pp. 3543–3559. https://doi.org/10.1109/AERO.2002.1035332.

[26] Črepinšek, M., Liu, S.-H., and Mernik, M., "Exploration and Exploitation in Evolutionary Algorithms: A Survey," *ACM Computing Surveys*, 2013. https://doi.org/10.1145/2480741.2480752.

[27] Preskill, J., "Quantum Computing in the NISQ Era and Beyond," *Quantum*, 2018. https://doi.org/10.22331/q-2018-08-06-79.

[28] Bharti, K., Cervera-Lierta, A., et al., "Noisy Intermediate-Scale Quantum Algorithms," *Reviews of Modern Physics*, 2022. https://doi.org/10.1103/RevModPhys.94.015004.