

Article

A UAV Path Planning Framework for LEO Satellite Monitoring with Hard Corridor Constraints

Zhixiang Li ^{1,2}, Desong Jiang ¹, Yu Hu ³, Haoran Li ⁴, Li Luo ¹, Ziqiao Tang ^{1,*}, Haiyin Qing ^{2,*} and Tao Liu ^{5,*} 

¹ School of Electronic Information and Electrical Engineering, Chengdu University, Chengdu 610106, China; lizhixiang@cdu.edu.cn (Z.L.); jiangdesong@cdu.edu.cn (D.J.); luoli_2021@163.com (L.L.)

² School of Electronic Information and Artificial Intelligence, Leshan Normal University, Leshan 614000, China

³ National Institute of Measurement and Testing Technology, Chengdu 610021, China; huyu@nimtt.com

⁴ School of Information and Control Engineering, Southwest University of Science and Technology, Mianyang 621010, China; horn_lee@swust.edu.cn

⁵ Chengdu Kinyea Technologies Co., Ltd., Chengdu 610000, China

* Correspondence: tangziqiao@cdu.edu.cn (Z.T.); qinghaiyin123@163.com (H.Q.); tovi.liu@kinyea.com (T.L.)

Abstract

Ground-based radar monitoring of low Earth orbit (LEO) satellites is constrained by terrain occlusion and short visibility windows. Deploying UAVs as mobile platforms can extend coverage, but beam alignment, obstacle avoidance, and kinematic limits have not been jointly addressed. This paper proposes a UAV path planning framework that treats the ground station–UAV–satellite (G–U–S) collinear relationship as a planning prior. Ideal UAV positions are first computed from satellite visible arcs. In obstacle-free airspace, these positions directly define the flight path. When no-fly zones intersect the ideal trajectory, a geometry-guided RRT* planner generates collision-free detours by favoring sampling near the beam axis. The beam safety corridor is then embedded with kinematic limits as a conservative hard constraint within a Minimum Snap Bézier QP model, whose convex hull property bounds the trajectory within the coverage region. Simulations over a representative visible arc of the target satellite (approximately 3.3 min) achieve 100% beam coverage in the obstacle-free scenario and 95.2% with three no-fly zones placed along the ideal trajectory. Visible arc duration varies with satellite orbital trajectory; the framework is applicable to any arc whose geometry satisfies the kinematic constraints of the UAV platform. An ablation experiment confirms that the geometry-guided sampling strategy is essential: replacing it with uniform random sampling reduces coverage by 9.5 percentage points and triggers a diagnostic fallback mechanism. Edge case stress testing further identifies the conditions under which beam containment and obstacle avoidance come into tension. The proposed framework provides a simulation-validated trajectory planning approach for UAV-assisted LEO satellite monitoring in complex airspace.

Keywords: UAV path planning; LEO satellite tracking; RRT*; Bézier trajectory optimization; hard corridor constraints; beam coverage



Academic Editor: Fabio D'Andreagiovanni

Received: 17 June 2026

Revised: 19 July 2026

Accepted: 20 July 2026

Published: 22 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

A single LEO satellite is visible to a fixed ground-based radar station for only a few minutes per pass, a limitation imposed by terrain occlusion and the curvature of the Earth. The deployment of unmanned aerial vehicles (UAVs) as mobile relay platforms has therefore been proposed to extend the usable observation window. In this scheme, the UAV is positioned along the line of sight between the ground station and the satellite

during each transit. In this relay configuration, a conical beam is transmitted by the ground radar toward the satellite, and the signal gain is maximized at the beam center, which lies on the line of sight from the ground station to the satellite. To maintain reliable signal reception, the UAV must stay within the narrow footprint of the beam at its flight altitude. This footprint defines the beam safety corridor, which tightly constrains the UAV flight trajectory. UAV–satellite cooperative monitoring falls within the space–air–ground integrated network paradigm [1]. Given their extended endurance and higher cruise speeds, fixed-wing UAVs are a practical choice for this relay role. Three stringent requirements are imposed on the flight trajectory: no-fly zones in low-altitude airspace must be avoided, the radar beam constraint must be satisfied, and the kinematic limits of the fixed-wing platform must be respected. When no-fly zones lie on the beam axis, obstacle avoidance forces the UAV away from the corridor, and the kinematic limits of the fixed-wing platform prevent instantaneous recovery of alignment. A direct conflict among these three requirements is therefore created. This combination of challenges has not been addressed jointly in existing path planning methods.

Existing path planning algorithms such as A*, Dijkstra, and RRT* [2] are primarily designed for path length minimization and collision avoidance. Communication-aware UAV path planning methods focus on throughput or coverage objectives, and the beam alignment condition, when included at all, is treated as a soft penalty term rather than a strict constraint. As a result, alignment accuracy may be traded off against other objectives, and strict collinearity cannot be guaranteed when the penalty weight and obstacle avoidance forces become comparable. The dynamic geometric alignment required for high-precision beam tracking has not been enforced as a hard boundary in any prior path planning framework.

This paper presents a trajectory planning framework that treats the ground station–UAV–satellite (G-U-S) collinear geometric constraint as a planning prior. Ideal UAV positions computed from the three-point collinear condition serve as geometric guidance throughout the pipeline. A dual-mode strategy selects between direct ideal-position tracking and geometry-guided RRT* depending on obstacle presence. The beam safety corridor is embedded as a conservative hard constraint within a Minimum Snap Bézier quadratic program (Section 4.4.3).

The main contributions of this paper are as follows:

- A trajectory planning framework is proposed that converts the G-U-S collinear geometric constraint from a post hoc coverage verification metric into a planning-stage guidance prior. By embedding the beam corridor as a conservative hard constraint within convex quadratic programming, the framework addresses the conflict among beam alignment, obstacle avoidance, and kinematic feasibility within a single planning framework.
- A geometric computation method for determining the ideal UAV position is introduced. This method locates the intersection of the ground station–satellite ray with the UAV altitude plane by solving the three-point collinear condition via numerical root-finding, whose unique solution is guaranteed by the intermediate value theorem.
- A geometry-guided RRT* algorithm is developed that combines goal-biased and corridor-constrained sampling to improve search efficiency within the narrow beam corridor.

The proposed framework differs from prior RRT*-based methods in three respects. First, whereas Informed RRT* biases sampling toward a path-length ellipsoid, the geometry-guided sampling introduced here steers the search toward the G-U-S beam axis, converting a geometric constraint into a search prior. Second, the sampling strategy is co-designed with the downstream QP: the biased samples supply waypoints within the convergence basin

of the corridor inequality, a coupling absent from planners that separate path planning from trajectory smoothing. Third, the diagnostic fallback provides a trajectory-availability guarantee with no counterpart in Informed RRT*, Dubins, or A*-QP methods. The ablation experiment in Section 5.4 confirms that removing the geometry guidance reduces coverage by 9.5 percentage points, demonstrating that each component is necessary.

The remainder of the paper is organized as follows. Related work is reviewed in Section 2. The system model and constraints are formulated in Section 3. The proposed algorithm is described in detail in Section 4. Simulation results and analysis are reported in Section 5, including an ablation experiment (Section 5.4), edge case stress testing (Section 5.5), and a discussion of parameter sensitivity and failure modes (Section 5.6). Conclusions are drawn in Section 6.

2. Related Work

This paper addresses the intersection of three research streams: LEO satellite orbit propagation and visibility prediction, UAV path planning under kinematic constraints, and trajectory optimization with geometric corridor constraints. A fourth stream, UAV-LEO satellite cooperation and beam alignment, provides the application context.

Ground-based LEO satellite monitoring depends on orbit prediction and visibility window estimation. Vallado et al. [3] developed the SGP4/SDP4 model from two-line element (TLE) data, and Kelso [4] later validated its accuracy against GPS precise ephemerides. Montenbruck and Gill [5] established the theoretical framework for computing ground station azimuth, elevation, and slant range; updated implementations appear in Vallado's fifth edition [6]. Recent work has documented limitations of the standard SGP4 pipeline. Kassas et al. [7] showed that closed-loop machine-learning correction improves LEO orbit prediction accuracy by an order of magnitude relative to raw SGP4 propagation. Kralfallah et al. [8] optimized ground tracking station deployment and found that approximately 95 stations are required for global triple-satellite coverage, which quantifies the scalability limit of fixed infrastructure. These works assume a stationary ground station constrained by terrain occlusion and Earth curvature. To overcome this limitation, researchers have introduced UAVs as mobile monitoring platforms, and trajectory design has become a central research topic.

UAVs have been extensively investigated as mobile platforms for extending ground monitoring capability. Ghambari et al. [9] published a comprehensive survey that classified UAV path planning techniques into classical, soft-computing, and hybrid strategies, covering developments through 2024. Zeng et al. [10] first studied trajectory optimization in UAV-assisted wireless communication and proposed a throughput-maximizing flight path design. The same group later incorporated fixed-wing kinematic constraints into an energy-efficient trajectory scheme [11]. Wu et al. [12] developed a joint trajectory and resource allocation framework for multi-UAV cooperation using successive convex approximation to reduce energy consumption. Li et al. [13] provided a comprehensive review of UAV communication for 5G and beyond. Although these works established the analytical foundation for UAV trajectory optimization, their objectives remain communication-centric. Throughput, energy efficiency, and service quality dominate the cost function, whereas geometric alignment with a ground radar beam is either absent or treated as a secondary factor.

The path planning problem beneath these communication-driven approaches shares algorithmic foundations with robotic motion planning. Karaman and Frazzoli [2] proposed the asymptotically optimal RRT* algorithm by adding a rewiring mechanism to standard RRT. Goal-biased variants subsequently accelerated convergence, and Informed RRT* [14] restricted sampling to a dynamically shrinking ellipsoid for faster optimal-path convergence. Recent improvements further enhance this family. Wu et al. [15] fused

Informed RRT* with adaptive probability-weighted sampling and a dynamic window approach for real-time obstacle avoidance. Fan et al. [16] combined bidirectional RRT* with artificial potential fields and a goal-biased strategy to reduce path length in complex environments. However, these RRT* variants target multi-rotor platforms and do not account for the non-holonomic constraints of fixed-wing UAVs. For fixed-wing platforms, Dubins-based planners capture minimum turning radius constraints [17], and Bézier curves combined with metaheuristic search satisfy bank-turn limits [18]. Burzyński and Stecz [19] incorporated fixed-wing dynamics into a multiplatform spacetime RRT* framework for multi-UAV swarm trajectory planning. Motamedi et al. [20] combined Dubins curves with ant colony optimization for minimum-time search across multiple fixed-wing UAVs under both static and dynamic obstacles. These methods advance fixed-wing kinematic constraint handling, yet none of them incorporate beam alignment constraints from ground-based radar tracking.

After a discrete reference path is obtained, a continuous and dynamically feasible trajectory must be generated. Mellinger and Kumar [21] proposed the Minimum Snap polynomial method for trajectory smoothing. Richter et al. [22] later introduced a Bézier representation to improve computational efficiency. Zhang et al. [23] formally analyzed the convex hull property of Bézier curves and showed that trajectory containment within a spatial corridor follows from linear inequality constraints on the control points alone. The FASTER algorithm [24] extended this approach by combining safe flight corridors with real-time optimization for quadrotors. Arshad et al. [25] formulated trajectory generation as a quadratic program with Obstacle-Free Corridor constraints and demonstrated convergence within approximately ten solver iterations. Sun et al. [26] proposed a two-layer framework that combines enhanced A* with quadratic programming, where the QP stage enforces corridor bounds and dynamic limits as hard constraints. These trajectory optimization and corridor-constraint methods were designed for terrestrial or indoor environments with static geometric corridors. Beam alignment constraints from ground-based radar tracking define a dynamic, satellite-following corridor whose position and width vary continuously with the satellite elevation angle. Such constraints have not been incorporated into any of these frameworks.

A separate body of work addresses UAV and LEO satellite cooperation, although with objectives distinct from radar beam alignment. Zhang et al. [27] studied resource allocation and offloading strategies while treating the flight path as a fixed input. Lee et al. [28] surveyed UAV-satellite communication architectures with an emphasis on protocol design. Mohamed et al. [29] formulated joint trajectory planning and LEO satellite selection as a multi-armed bandit problem. Chai et al. [30] constructed a resource scheduling system for UAV and LEO edge computing. Amponis et al. [31] proposed a composite cost metric for energy-efficient routing in UAV-satellite non-terrestrial networks. Mohapatra [32] introduced an adaptive ant colony optimization framework for dynamic task allocation and relay scheduling across UAV swarms and LEO satellite constellations. A small number of recent studies have started to address the beam alignment problem. Cui et al. [33] developed sensing-assisted mmWave beam management that uses extended Kalman filtering to predict beam directions from visual-domain features and UAV trajectory estimates. Zhao et al. [34] addressed beam tracking for UAV-mounted satellite communication on-the-move with massive antenna arrays, treating beam alignment as a signal-processing rather than a path-planning problem. Across all these works, the UAV trajectory is optimized for communication throughput, resource efficiency, or beamforming gain. The geometric relationship with the ground radar beam is either absent or treated as a secondary factor. When beam alignment is considered, it appears as a soft penalty term whose satisfaction cannot be guaranteed under competing objectives. No existing path planning method

enforces, as a hard geometric constraint embedded in the feasible domain of trajectory optimization, the requirement that the ground station, the UAV, and the satellite remain collinear within a narrow angular tolerance.

Recent work in non-terrestrial network (NTN) and satellite IoT trajectory design further underscores this gap. Lukito et al. [35] integrated STAR-RIS with UAVs for LEO satellite IoT and jointly optimized flight path, RIS phase shifts, and power allocation, achieving 67% higher sum rates. Tesfaw et al. [36] proposed a federated multi-agent deep reinforcement learning framework for joint 3D UAV trajectory and resource allocation in LEO satellite-enabled marine IoT networks, demonstrating superior energy efficiency over single-agent baselines. Yao et al. [37] designed RIS-UAV-assisted LEO satellite constellation communication with alternating optimization of satellite beamforming, RIS configuration, and UAV trajectory. Yu et al. [38] addressed joint 3D beamforming and trajectory design for UAV-satellite uplink covert communication, explicitly coupling beam pattern geometry with flight path. Liang et al. [39] investigated satellite-assisted UAV control with joint sensing and communication scheduling for IoT data collection, treating trajectory tracking as a constrained optimization problem. He et al. [40] studied NOMA-based UAV-assisted satellite IoT with joint access, power, and trajectory optimization via Taylor convex approximation. Across these NTN-oriented studies, the optimization objective is consistently communication-centric, targeting throughput maximization, energy efficiency, or covert-ness, while the geometric beam alignment required for radar-based satellite monitoring is not addressed. The beamforming or RIS phase configuration is optimized for signal quality at the receiver, not for maintaining a rigid collinear geometry among ground station, relay, and space asset.

Three gaps summarize the current literature. First, existing communication-aware and path planning methods do not incorporate the dynamic collinear G-U-S geometric constraint. The beam alignment requirement is either omitted or relegated to a post hoc verification metric. Second, no approach jointly addresses radar beam alignment, obstacle avoidance, and fixed-wing kinematic constraints within a single planning framework. When any two of these three requirements are satisfied, the third is typically relaxed or ignored. Third, beam alignment has been handled exclusively through soft penalties, and their stability under competing constraints is not guaranteed. The method presented in this paper closes these three gaps. It converts the G-U-S collinearity condition from a verification metric into a planning prior. It embeds the beam safety corridor as a hard inequality constraint within a convex quadratic programming formulation. It resolves the conflict between beam alignment and obstacle avoidance through geometry-guided biased sampling in the RRT* stage.

3. System Modeling and Constraints

3.1. Problem Description

In the simulation scenario, a fixed-wing UAV is employed as a mobile observation platform operating at a fixed altitude. During each visible satellite arc, the UAV is required to satisfy the collinear geometric constraint among the ground station, the UAV, and the satellite as closely as possible, thereby extending the effective observation duration. No-fly zones and other obstacles may be present in the low-altitude environment. The proposed method must minimize the UAV's deviation from the beam axis while ensuring obstacle avoidance, so that obstacle avoidance and beam alignment are jointly satisfied. In addition, the kinematic constraints of the fixed-wing UAV, including velocity and acceleration limits, must be satisfied, and a flyable trajectory must be produced.

To achieve these objectives, two core problems must be solved. The first is to determine, given the satellite TLE data, the ground station location, the radar beam half-cone angle, and

the minimum elevation threshold, the ideal UAV position satisfying the G-U-S collinearity constraint at each time step on the assigned fixed altitude. The second is to generate, when no-fly zones exist in the airspace, a smooth, dynamically feasible trajectory that remains as close as possible to the ideal position while satisfying obstacle avoidance and maintaining the fixed altitude.

To formalize these problems quantitatively, the relevant coordinate systems are first defined in Section 3.2, and the model construction is then presented in Section 3.3.

3.2. Coordinate System Definition and Transformation

To accurately describe the dynamic geometric relationship among the ground station, the UAV, and the satellite, four coordinate frames are used throughout the modeling: Earth-centered inertial (ECI) for satellite orbital motion, Earth-centered Earth-fixed (ECEF) for ground station position, WGS-84 geodetic (LLH) for geographic coordinates, and East–North–Up (ENU), a local Cartesian frame centered at the ground station whose Up component defines the UAV fixed altitude.

The purpose of defining and transforming these coordinate systems is to unify the coordinates of the satellite, the ground station, and the UAV into the ENU frame. With this unified framework, the ideal UAV position under the collinearity constraint can be derived, path planning can be performed in the ENU frame, and the azimuth, elevation, and slant range required for beam alignment can be computed. Based on the above coordinate systems, three coordinate transformation workflows are established, as shown in Figure 1.




Figure 1. Coordinate transformation workflows for the satellite, ground station, and UAV.

Subsequently, both the ECEF vector from the ground station to the satellite and the UAV position vector are transformed to the ENU frame via rotation matrices. The azimuth A , elevation ε , and slant range R are obtained from the ENU coordinates according to Equation (1):

$$\begin{aligned}
 A &= \arctan(E, N), \\
 \varepsilon &= \arcsin(U/R), \\
 R &= \sqrt{E^2 + N^2 + U^2},
 \end{aligned} \tag{1}$$

where azimuth, elevation, and slant range describe the spatial orientation and distance of a target relative to the ground station, by which the three-point collinearity constraint modeling in Section 3.3 is supported.

3.3. Model Construction

3.3.1. Three-Point Collinearity Model

The three-point collinearity model establishes the geometric condition from which the ideal UAV position is determined at each time step. Precise beam alignment is fundamental to maintaining sufficient signal strength for reliable downlink reception [34]. To maintain reliable signal reception, the UAV must be kept within the effective beam coverage of the ground-based radar, as any deviation in beam alignment may lead to a non-negligible gain loss. For this reason, the geometric constraint of three-point collinearity among the ground

station, the UAV, and the satellite is imposed as the ideal condition. Because the UAV is operated at a fixed altitude, its ideal position at any time instant is uniquely determined by the azimuth and elevation of the line-of-sight vector from the ground station to the satellite in the ENU frame, i.e., by the intersection of that ray with the horizontal plane at the specified altitude.

To compute the ideal UAV position quantitatively, the azimuth A and elevation ε of the satellite relative to the ground station are first obtained from Equation (1). Together these two angles define the unit direction vector of the ray from the ground station toward the satellite. The ray is then parameterized in the ECEF frame and expressed as:

$$X(t) = X_G + t \cdot \Delta X_{GS}, \quad t \in [0, 1] \quad (2)$$

where $X_G = (X_0, Y_0, Z_0)$ denotes the position vector of the ground station in the ECEF frame, and (X_s, Y_s, Z_s) are the ECEF coordinates of the satellite obtained from the SGP4 propagation. The vector $\Delta X_{GS} = (X_s - X_0, Y_s - Y_0, Z_s - Z_0)$ is the ECEF coordinate difference from the ground station to the satellite. The scalar parameter t sweeps the ray continuously from the ground station, at $t = 0$, to the satellite, at $t = 1$.

The ideal UAV position is defined as the point on the ray whose height equals the prescribed fixed flight altitude H_{target} . This condition leads to the scalar equation:

$$f(t) = H(t) - H_{\text{target}} \quad (3)$$

where $H(t)$ is the ellipsoidal height of the point $X(t)$ obtained by transforming $X(t)$ from the ECEF frame to the WGS-84 geodetic coordinate system (LLH) through the standard conversion procedure described in Section 3.2, and H_{target} is the prescribed fixed flight altitude of the UAV.

A unique solution to Equation (3) in $t \in [0, 1]$ is guaranteed by the following geometric argument. At $t = 0$ the ray lies at the ground station, whose ellipsoidal height is near zero and therefore well below $H_{\text{target}} = 1500$ m. At $t = 1$ the ray reaches the satellite at an orbital altitude of approximately 550 km, far above H_{target} . Because $H(t)$ varies continuously along the ray segment, the intermediate value theorem ensures the existence of at least one root t^* for which $H(t^*) = H_{\text{target}}$. Furthermore, as the ray proceeds monotonically upward along the segment throughout the visible arc, this root is unique. The root t^* is obtained via the MATLAB R2024a `fzero` function with a convergence tolerance of 10^{-6} on t , corresponding to a height residual below 1 m. Once t^* is substituted back into Equation (2), the ECEF coordinates of the ideal UAV position are recovered and subsequently converted to the geodetic and ENU frames for use in the path planning stage.

3.3.2. System Constraints

In this study, a fixed-wing UAV is employed as the mobile observation platform. The following kinematic constraints are imposed to ensure flight safety and trajectory feasibility.

1. Kinematic Constraints

(1) Speed constraint

The instantaneous speed $V_{\text{UAV}}(t)$ of the UAV is constrained not to exceed the maximum allowable speed V_{max} :

$$\|V_{\text{UAV}}(t)\| \leq V_{\text{max}} \quad (4)$$

In this simulation, V_{max} is set to 50 m/s.

(2) Acceleration constraint

The acceleration of the UAV is limited by the maximum allowable acceleration:

$$\|a_{\text{UAV}}(t)\| \leq a_{\text{max}} \quad (5)$$

In this simulation, the maximum acceleration is set to 5 m/s^2 to ensure trajectory smoothness.

(3) Altitude constraint

The safe corridor radius, defined as the projection of the radar beam cross-section onto the UAV altitude plane, is given for a beam half-cone angle $\theta = 2^\circ$ at satellite elevation ε by:

$$R_c = \frac{H \cdot \tan \theta}{\sin \varepsilon} \quad (6)$$

where R_c denotes the safe corridor radius and H is the fixed UAV altitude.

The required corridor radius is set by the expected lateral deviation during obstacle avoidance. Three no-fly zones of maximum radius 55 m are centered on the ideal trajectory (Section 4). The RRT* detour imposes a lateral deviation of at least one obstacle radius ($\sim 55 \text{ m}$), and the 20 m extension step adds discretization clearance. With a 20% contingency margin for trajectory smoothing and transient effects, the required corridor radius is approximately 90 m. Substituting $R_c = 90 \text{ m}$ and $\varepsilon = 30^\circ$ into Equation (6) yields a minimum altitude $H_{\text{min}} \approx 1289 \text{ m}$.

Table 1 compares candidate altitudes. The baseline value of 716 m, which is obtained from Equation (6) with $R_c = 50 \text{ m}$ under the assumption of no obstacle avoidance requirement. At this altitude, the corridor radius ranges from 25.0 to 50.0 m, which is narrower than the largest obstacle (55 m) at all elevations above 45° . At 1000 m and 1200 m, the corridor remains insufficient at the lowest elevation. At 1500 m, the minimum radius reaches 104.7 m, providing 14.7 m of margin above the required 90 m at $\varepsilon = 30^\circ$. Further increasing the altitude to 1800 m increases the flight distance to 251% of the baseline without proportional benefit. Accordingly, $H = 1500 \text{ m}$ is adopted for all experiments.

Table 1. Comparison of UAV altitude candidates.

$H \text{ (m)}$	$R_c \text{ at } \varepsilon = 30^\circ \text{ (m)}$	$R_c \text{ at } \varepsilon = 90^\circ \text{ (m)}$	Margin Above 90 m (m)	Flight Distance Ratio
716	50.0	25.0	−40.0	1.00
1000	69.8	34.9	−20.2	1.40
1200	83.8	41.9	−6.2	1.68
1500	104.7	52.4	+14.7	2.09
1800	125.7	62.8	+35.7	2.52

Flight distance ratio is defined as the total path length at altitude H relative to that at the baseline $H = 716 \text{ m}$.

2. Beam Geometric Constraints

The ground-based radar transmits a conical beam with half-cone angle $\theta = 2^\circ$ toward the satellite. Let $U_{\text{sat}}(t)$ and $U_{\text{uav}}(t)$ denote the unit vectors from the ground station to the satellite and to the UAV, respectively. The UAV lies within the beam when:

$$U_{\text{sat}}(t) \cdot U_{\text{uav}}(t) \geq \cos \theta \quad (7)$$

Rather than requiring exact collinearity, the proposed framework allows the UAV to deviate laterally from the beam axis, provided the offset remains within the safe corridor radius $R_c(t)$ from Equation (6):

$$\|p(t) - p_{\text{ideal}}(t)\| \leq R_c(t), \quad \forall t \quad (8)$$

where $p(t)$ is the instantaneous UAV position in the ENU frame and $p_{\text{ideal}}(t)$ is the ideal position from Section 3.3.1.

3. Constraint Classification

Based on the severity of their violation consequences, the constraints introduced above are categorized into two classes:

(1) Hard constraints

The beam corridor constraint (Equation (8)) and the kinematic limits on speed and acceleration (Equations (4) and (5)) bear directly on mission success and flight safety and are treated as non-negotiable. The beam corridor is embedded as a linear inequality within the QP feasible domain (Section 4); the kinematic constraints are enforced through distance-proportional segment time allocation and verified after optimization.

(2) Soft constraint

Trajectory smoothness, quantified by the snap energy, tolerates moderate degradation without causing mission failure. Smoothness is therefore incorporated as a penalty term in the objective function, allowing the optimizer to balance smoothness against the hard constraints.

3.3.3. System Scenario Overview

Figures 2 and 3 illustrate the two scenarios addressed in this paper. Figure 2 shows the obstacle-free case, where the UAV follows the beam axis without deviation. Figure 3 introduces no-fly zones intersecting the flight path, requiring a detour within the safety corridor.




Figure 2. Obstacle-free scenario. The UAV remains within the radar beam cone directed from the ground station toward the satellite, by which the collinear G-U-S geometry is maintained. The trajectory is bounded by the safety corridor defined by the 2° half-cone beam angle.




Figure 3. Scenario with no-fly zones. (a) Side-view system schematic, in which cylindrical restricted regions intersect the UAV flight path. (b) Top-down view of the detour path circumventing the circular no-fly zones, where the dashed circle denotes the safety corridor boundary.

4. Algorithm

4.1. Algorithm Framework

The proposed method follows a four-stage unified framework, namely visible arc extraction, dual-mode path planning, trajectory smoothing optimization, and coverage assessment, as illustrated in Figure 4. The data flow and core functionality of each stage are described below.




Figure 4. Proposed algorithm framework. The G-U-S collinearity constraint is integrated at three levels : (i) ideal position computation in Stage 1, (ii) geometry-biased sampling in Stage 2, and (iii) a conservative corridor inequality embedded in the Stage 3 Bézier QP. The pipeline consists of four stages: (1) visible arc extraction and ideal position generation, (2) dual-mode path planning, (3) Minimum Snap Bézier trajectory optimization with conservative corridor constraint and diagnostic fallback, and (4) beam coverage assessment and trajectory export.

Stage 1: The satellite azimuth-elevation-range (AER) time series is computed using the SGP4 orbital propagation model. Continuous visible arcs are extracted by applying a minimum elevation threshold, which establishes the temporal skeleton for subsequent planning. For each visible time step, the ideal UAV position is determined by solving the three-point collinearity condition at the designated fixed altitude. The resulting ideal position sequence is employed as a unified geometric guidance reference for both path planning modes.

Stage 2: Path planning is performed independently for each visible arc. The ideal positions are first transformed into a local East–North–Up (ENU) coordinate frame with the ground station as the origin. The radius of the safe corridor is determined from the beam half-cone angle and the instantaneous satellite elevation. No-fly zones spatially intersecting the current segment are then identified, and the planning mode is selected automatically based on the presence or absence of obstacles. In the obstacle-free case, the ideal position sequence is adopted directly as the reference path, as detailed in Section 4.2. In the obstacle case, a geometry-guided RRT* planner incorporating goal biasing and corridor-constrained sampling is employed to generate a collision-free detour path, as detailed in Section 4.3. The resulting path is subsequently mapped to uniform time steps via arc-length parameterization. Both modes produce discrete reference waypoints of a unified format, which are transferred to Stage 3.

Stage 3: The reference waypoints from Stage 2 are received, and the flight time of each segment is allocated in proportion to its path length. The trajectory within each segment is parameterized as a Bézier curve. A weighted objective function combining a snap minimization term and a deviation penalty term is formulated. Endpoint fixation and inter-segment continuity are imposed as equality constraints, while the safe corridor boundaries are enforced as inequality constraints within a convex quadratic programming (QP) formulation. As noted in Section 3.3.2, the feasibility of this QP depends on the quality of the reference path supplied by Stage 2: when the waypoints remain close to the beam axis, the corridor bounds define a non-empty feasible domain and the QP converges reliably. By virtue of the convex hull property of the Bézier representation, constraining the control points to lie within the corridor bounds the trajectory to the coverage region. As analyzed in Section 4.4.3, this constraint is a conservative linear approximation of the circular beam cross-section; continuous-time containment is assessed at discrete sampling instants after optimization. After the QP is solved, a diagnostic check is performed. If the fraction of trajectory points that deviate beyond the beam boundary exceeds a preset threshold, a fallback strategy is triggered. In this strategy, the reference path is retained and kinematic quantities are estimated via finite differences, which ensures trajectory availability under all conditions. The complete optimization procedure is detailed in Section 4.4.

Stage 4: The smoothed trajectory is converted to the angular domain. At each time step, the spherical angular distance between the UAV and satellite line-of-sight vectors is computed. Beam coverage status is determined by verifying that this angular distance falls within the half-cone angle of the radar beam. The overall coverage rate is computed, and the velocity and acceleration profiles are verified against the prescribed kinematic limits throughout the entire observation window.

The framework runs offline: satellite ephemeris is propagated from TLE data via the SGP4 model before the mission, ideal UAV positions are computed from the G-U-S collinear condition, and the full trajectory is generated before takeoff. This pre-mission strategy exploits two properties of the target domain. LEO satellite orbits are deterministic for several hours ahead, and no-fly zones in low-altitude airspace are typically static or evolve on timescales far longer than a single visible arc. We identify only two scenarios that would trigger online replanning: an updated TLE data set becoming available mid-mission, which

alters the predicted satellite track, or a previously unknown obstacle appearing on the planned path and detected by onboard sensors. Both extensions are discussed in Section 6.

We treat the G-U-S collinearity constraint as a planning-stage prior, not a post hoc check. The same geometric condition threads through three pipeline stages. Stage 1 computes ideal positions from three-point collinearity. Stage 2 biases RRT* sampling toward the beam axis using those positions as guidance. Stage 3 embeds the beam corridor as a conservative inequality within the QP feasible domain. The cumulative effect is that each stage delivers its output into the convergence basin of the next: the RRT* planner receives target waypoints that lie near the beam axis, and the QP solver receives waypoints that already satisfy the corridor constraint to within the sampling resolution. If the geometric prior were applied only at the QP stage, with no guidance at the sampling stage, RRT* waypoints could deviate far from the axis, shrinking the QP feasible domain to the point of infeasibility. The ablation experiment (Section 5.4) confirms this prediction: removing the sampling-stage guidance reduces coverage by 9.5 percentage points and triggers the diagnostic fallback.

4.2. Obstacle-Free Path Planning

In the obstacle-free scenario, the ideal collinear trajectory is not intersected by any no-fly zones. The ideal position sequence computed in Stage 1 already satisfies the G-U-S collinear constraint at every time step. This sequence is therefore adopted directly as the reference path for trajectory smoothing, and no additional path planning is required.

Since the reference waypoints are already expressed in the ENU coordinate frame, the flight time of each segment between consecutive waypoints is allocated in proportion to the Euclidean distance of that segment. The waypoints and segment times are then transferred to the Minimum Snap Bézier optimizer (Section 4.4), which generates a continuous, dynamically feasible trajectory that remains within the safe corridor. The Bézier order n_b is set to $\min(7, \max(3, N_{wp} - 1))$, where N_{wp} denotes the number of waypoints in the current visible arc, to ensure at least cubic smoothness while limiting the number of decision variables. The reference path inherently lies along the beam center axis, and only minor adjustments are required of the optimizer to enforce smoothness and kinematic limits. The resulting trajectory is therefore expected to achieve full beam coverage.

The complete procedure is summarized in Algorithm 1.

Algorithm 1 Obstacle-Free Path Planning

Input: Ideal waypoints, time vector, corridor safety radius

Output: Smoothed trajectory, velocity, acceleration

1. reference path = Ideal waypoints
 2. N_{wp} = number of waypoints
 3. segment lengths = consecutive Euclidean distances between waypoints
 4. total length = sum(segment lengths)
 5. **if** total length $< 10^{-6}$ **then**
 6. segment times = (total segment time / ($N_{wp} - 1$)) for all segments
 7. **else**
 8. segment times = (segment lengths / total length) $\times$ total segment time
 9. **end if**
 10. $n_b = \min(7, \max(3, N_{wp} - 1))$
 11. corridor lower bound = reference path – corridor radius
 12. corridor upper bound = reference path + corridor radius
 13. Solve Minimum Snap Bezier QP with conservative corridor constraint (C^1 continuity, corridor bounds) to obtain control points
-

Algorithm 1 *Cont.*

```

14. if QP does not converge then
15.   Use linear interpolation initial guess as control points
16. end if
17. Evaluate Bezier trajectory at query times to obtain pos, vel, acc
18. deviation = Euclidean distance between position and ideal waypoints
19. exceedance ratio = count(deviation > corridor radius)/length(deviation)
20. if exceedance ratio > 0.2 then
21.   Retain reference path, recompute vel and acc via finite differences
22. end if
23. return position, velocity, acceleration

```

4.3. Obstacle Mode: Geometry-Guided RRT* Path Planning

When no-fly zones intersect the ideal collinear trajectory, the obstacle mode is activated in Stage 2. In this mode, the ideal position sequence is employed as the geometric guidance reference, and a geometry-guided RRT* algorithm is used to generate a collision-free detour path that remains as close as possible to the beam center axis.

The algorithm is executed independently between each pair of consecutive ideal waypoints. If the direct line segment connecting two adjacent waypoints is found to intersect an obstacle, a local RRT* search is triggered. For each obstructed segment, a random tree is grown within a bounded search region, and standard RRT* operations are performed in sequence: nearest-neighbor selection, constrained extension, parent re-selection, and rewiring. The near-neighbor search radius is gradually reduced as the number of nodes in the tree increases, preserving the asymptotic optimality of the algorithm. The search for a given segment is terminated early when the distance between the new node and the local goal is reduced below 1.5 times the extension step size.

In standard RRT*, the sampling distribution is uniform over the configuration space. Convergence is improved by informed variants such as Informed RRT*, in which samples are restricted to an ellipsoidal region whose focal points are the start and goal and whose shape contracts as better paths are found, thereby focusing exploration on regions that can improve the current solution cost [14]. These methods, however, are driven by Euclidean path length minimization and are not designed for mission-specific geometric constraints. In the satellite monitoring context, path quality is not measured by distance alone but by adherence to the beam corridor. A short detour path that stays within the corridor is preferable to a minimum-length path that exits it. The cost-driven ellipsoidal heuristic is therefore replaced by a beam-aligned corridor derived from the G-U-S collinear geometry, which guides the random tree toward a sequence of ideal waypoints rather than a single destination.

To improve search efficiency in the narrow beam corridor, a biased sampling strategy is introduced. With probability p_{goal} , the local goal waypoint is directly used as the random sample, accelerating tree growth toward the target. With probability p_{geom} , the random sample is drawn within a rectangular region centered on the segment connecting the start and goal, with its half-width set to $4 \cdot \text{stepSize}$. This value is chosen so that the biased sampling region covers the expected lateral deviation of approximately 75 m derived in Section 3.3.2, since the extension step size is set to 20 m and four times this value yields 80 m. With the remaining probability $1 - p_{\text{goal}} - p_{\text{geom}}$, samples are drawn uniformly within the search bounds to retain probabilistic completeness. This biased sampling strategy substantially reduces the proportion of ineffective samples in the confined corridor. The values of p_{goal} and p_{geom} are held constant across all experiments and are listed in the parameter table (Table 2).

Table 2. Simulation parameters.

Parameter	Value	Description
Ground station position	30.5279° N, 114.888° E	Wuhan station
Satellite orbital altitude	~550 km	Near-circular LEO
Satellite orbital inclination	53°	NORAD 44717
UAV flight altitude	1500 m	Fixed altitude layer
Maximum flight speed	50 m/s	Fixed-wing platform constraint
Maximum acceleration	5 m/s ²	Kinematic constraint
Radar minimum elevation	30°	Visibility threshold
Beam half-cone angle	2°	G-U-S alignment constraint
Simulation time step	10 s	Trajectory sampling interval
TLE data scan window	12 h	Full period searched for visible arcs
Representative visible arc	3.3 min	Duration varies with satellite trajectory
Bezier optimization weights	$\lambda_{\text{snap}} = 1.0, \lambda_{\text{dev}} = 10.0$	Smoothness/Reference tracking
No-fly zone configuration	See Table 3	Three circular obstacles
RRT* goal bias p_{goal}	0.12	Probability of direct goal sampling
RRT* geometry bias p_{geom}	0.25	Probability of corridor-constrained sampling
RRT* extension step size	20 m	Maximum single-step tree extension
RRT* max iterations	2000	Per-segment sampling budget
RRT* goal threshold	$1.5 \times \text{stepSize}$	Early termination distance
QP fallback threshold	20%	Max. out-of-beam fraction before fallback

Table 3. No-fly zone parameters.

Obstacle No.	Relative Position	On Ideal Path	Radius (m)
1	0.28	Yes	45
2	0.52	Yes	55
3	0.74	Yes	40

After local planning is completed for all obstructed segments, the resulting path nodes are mapped to uniform time steps via arc-length parameterization. The discrete waypoint sequence is then transferred as the reference path to Stage 3 for trajectory smoothing optimization. The full procedure is summarized in Algorithm 2.

Algorithm 2 Geometry-Guided RRT***Input:** $p = \{p_1, \dots, p_M\}$, obstacles, bounds, stepSize, maxIter**Output:** path, n_{detours}

```

1. path =  $\{p_1\}$ ;  $n_{\text{detours}} = 0$ 
2. for  $i = 1$  to  $M - 1$  do
3.   if not Collision( $p_i, p_{i+1}$ , obstacles) then
4.     path = path  $\cup \{p_{i+1}\}$ 
5.     continue
6.   end if
7.    $n_{\text{detours}} = n_{\text{detours}} + 1$ 
8.    $V = \{p_i\}$ ; cost( $p_i$ ) = 0
9.   for iter = 1 to maxIter do
10.     $r = \text{Uniform}(0, 1)$ 
11.    if  $r < p_{\text{goal}}$  then
12.       $q_{\text{rand}} = p_{i+1}$ 
13.    else if  $r < p_{\text{goal}} + p_{\text{geom}}$  then
14.       $q_{\text{rand}} = (1 - t) p_i + t \cdot p_{i+1} + \varepsilon$ 
        where  $t \sim U(0, 1)$ ,  $\varepsilon \sim U(-4 \cdot \text{stepSize}, 4 \cdot \text{stepSize})$ 
15.    else
16.       $q_{\text{rand}} = \text{SampleFree}(\text{bounds}, \text{obstacles})$ 
17.    end if
18.     $q_{\text{nearest}} = \text{Nearest}(V, q_{\text{rand}})$ 
19.     $q_{\text{new}} = \text{Steer}(q_{\text{nearest}}, q_{\text{rand}}, \text{stepSize})$ 
20.    if Collision( $q_{\text{nearest}}, q_{\text{new}}$ , obstacles) then
21.      continue
22.    end if
23.     $V = V \cup \{q_{\text{new}}\}$ 
24.    Rewire( $V, q_{\text{new}}$ )
25.    if  $\|q_{\text{new}} - p_{i+1}\| \leq 1.5 \cdot \text{stepSize}$  then
26.      break
27.    end if
28.  end for
29.  if the goal is reached then
30.    path = path  $\cup \text{Backtrack}(q_{\text{new}})$ 
31.  else
32.    path = path  $\cup \text{BestPathFound}(V, p_{i+1})$ 
33.  end if
34. end for
35. return path,  $n_{\text{detours}}$ 

```

4.4. Minimum Snap Trajectory Optimization with Hard Constraints

The discrete waypoints produced by both modes are sampled position sequences for which velocity continuity between consecutive points is not guaranteed, and therefore cannot be directly executed by the fixed-wing UAV flight controller. To address this, the Minimum Snap Bézier curve optimization method is employed to smooth the discrete waypoints from both modes into a continuous trajectory under the safe corridor constraint, with the beam coverage corridor embedded as a conservative constraint in the optimization model.

4.4.1. Bézier Curve Parameterization and Time Allocation

To ensure high-order continuity of the flight trajectory, the path between each pair of consecutive waypoints is parameterized as an n_b -th order Bézier curve, where $n_b =$

$\min(7, \max(3, N_{wp} - 1))$ and N_{wp} is the number of waypoints in the current visible arc. The Bézier polynomial for the i -th segment is expressed as:

$$p_i(t) = \sum_{j=0}^{n_b} c_{i,j} B_j^{n_b} \left(\frac{t - t_{i-1}}{T_i} \right) \quad (9)$$

where $c_{i,j}$ denotes the j -th control point of the i -th segment, $B_j^{n_b}(\tau) = \binom{n_b}{j} \tau^j (1 - \tau)^{n_b-j}$ is the Bernstein basis function of degree n_b , and T_i is the flight time allocated to the segment. The segment times are assigned in proportion to the Euclidean distances of the corresponding path segments relative to the total path length.

Two geometric properties of Bézier curves are instrumental in this formulation. It is guaranteed by the convex hull property that the entire curve is contained within the convex hull of its control points; consequently, when all control points are confined to the safe corridor, the continuous trajectory is confined to that corridor. As discussed in Section 4.4.3, the corridor constraint is a conservative rectangular approximation of the circular beam cross-section. Exact passage of the curve through the first and last control points is ensured by the endpoint interpolation property, which permits the waypoints to serve as exact segment boundary anchors. By these two properties, the continuous corridor constraint is reduced to a finite set of linear inequalities on the control points, and the optimization problem is substantially simplified while strict feasibility is preserved.

4.4.2. Objective Function Formulation

The objective of trajectory optimization is that a balance between smoothness and tracking accuracy is achieved: the higher-order derivatives of the trajectory are minimized to ensure flight quality, while the optimized trajectory is kept as close as possible to the reference path determined by the G-U-S collinear geometry. To this end, the overall objective function is formulated as a weighted sum of a snap cost term and a deviation penalty term:

$$\min J = \lambda_{\text{snap}} J_{\text{snap}} + \lambda_{\text{dev}} J_{\text{dev}} \quad (10)$$

where J_{snap} is defined as the integral of the squared fourth derivative of the trajectory over all segments, by which trajectory smoothness is quantified:

$$J_{\text{snap}} = \sum_{i=1}^{N_{wp}-1} \int_0^{T_i} \|p_i^{(4)}(t)\|^2 dt \quad (11)$$

The deviation penalty term J_{dev} is defined as the quadratic penalty on the deviation of the control points from their reference positions:

$$J_{\text{dev}} = \sum_{i=1}^{N_{wp}-1} \sum_{j=1}^{n_b} \|c_{i,j} - p_{i,j}^*\|^2 \quad (12)$$

In the above expressions, $p_i^{(4)}(t)$ is the fourth-order time derivative of the i -th segment trajectory. The symbol $c_{i,j}$ denotes the j -th control point of the i -th segment. The symbol $p_{i,j}^*$ denotes the ideal reference position for that control point, obtained by linear interpolation along time between the consecutive G-U-S collinear waypoints.

The weights λ_{snap} and λ_{dev} govern the trade-off between smoothness and reference tracking. The former suppresses higher-order derivatives, while the latter attracts the trajectory toward the G-U-S reference path. Under the nominal values in Table 2, moderate deviation from the reference path is permitted as long as the corridor bounds are satisfied. Excessively large λ_{dev} can shrink the QP feasible domain, as analyzed in Section 5.6.2.

4.4.3. Hard Constraint Embedding and QP Solution

To enforce kinematic feasibility and embed the beam coverage requirement, the safe corridor constraint and the inter-segment continuity conditions defined in Sections 3.3.1 and 3.3.2 are incorporated directly into the optimization model as boundaries of the QP feasible domain.

1. Equality constraints

Three types of equality conditions are imposed to guarantee the connectivity and endpoint accuracy of the trajectory. First, the last control point of each segment is made to coincide with the first control point of the next segment, by which position continuity across segments is enforced. Second, the velocities of adjacent segments are matched at their junction, by which inter-segment velocity continuity (C^1 continuity) is enforced. Third, the first control point of the first segment and the last control point of the final segment are fixed to the given start and end waypoints, respectively. These conditions are collectively written as $A_{\text{eq}} c = b_{\text{eq}}$.

2. Inequality constraints

As discussed in Section 4.4.1, confining all control points within the safe corridor constrains, via the convex hull property, the entire trajectory to the coverage region. Each control point is therefore bounded by:

$$p_{i,j}^*(t_k) - R_c(t_k) \leq c_{i,j} \leq p_{i,j}^*(t_k) + R_c(t_k) \quad (13)$$

where $p_{i,j}^*$ is the reference position of the control point obtained by linear interpolation between the consecutive waypoints, and the corridor radius R_c is determined from the beam half-cone angle and the satellite elevation angle as defined in Equation (6).

Equation (13) imposes independent coordinate-wise bounds of half-width R_c on each control point, defining a rectangular feasible region that circumscribes the circular beam cross-section. Because a square of side $2R_c$ contains the inscribed circle of radius R_c , constraining the control points within the square does not mathematically guarantee that the trajectory lies within the circle: control points placed at the corners of the feasible rectangle could, in principle, place a portion of the Bézier curve outside the circular beam boundary. Two design choices render this conservatism benign in practice. First, the corridor radius R_c used in the QP is the minimum value over the visible arc (Section 4.2), whereas the actual beam radius at higher elevations is substantially larger (see Table 1). Second, the deviation penalty term J_{dev} in Equation (12) attracts the control points toward the beam axis, keeping them near the center of the feasible region rather than at its corners. Across all experiments reported in Sections 5.2–5.5, the maximum lateral deviation of the Bézier trajectory from the beam axis remained below $0.85R_c$, empirically within the circular beam boundary. The formulation is therefore a conservative hard constraint: a computationally tractable linear approximation of the circular beam constraint whose empirical containment performance matches that of the full circular formulation under the operating conditions considered.

This constraint is a set of linear inequalities, written as $A_{\text{ineq}} c \leq b_{\text{ineq}}$. A standard convex quadratic programming problem is constituted by the objective function and all constraints assembled above:

$$\min \frac{1}{2} c^T \mathbf{H} c + \mathbf{f}^T c \quad \text{s.t.} \quad A_{\text{eq}} c = b_{\text{eq}}, \quad A_{\text{ineq}} c \leq b_{\text{ineq}} \quad (14)$$

where $c \in \mathbb{R}^{n_c}$ is the decision vector concatenating all Bézier control point coordinates across all segments, $\mathbf{H} \in \mathbb{R}^{n_c \times n_c}$ is the Hessian matrix encoding the quadratic snap cost

and deviation penalty, and $\mathbf{f} \in \mathbb{R}^{n_c}$ is the linear coefficient vector arising from the deviation penalty term (the gradient of J_{dev} with respect to c). The QP is solved by a standard solver (MATLAB quadprog, interior-point algorithm). The velocity and acceleration constraints are implicitly satisfied through the distance-proportional segment time allocation described in Sections 4.2 and 4.3, and the peak acceleration is verified analytically via the Bezier derivative evaluation. The resulting control point vector c fully defines a smooth, kinematically feasible flight trajectory. Containment within the beam coverage region is assessed at discrete sampling instants as discussed above.

4.4.4. Diagnostic Check and Fallback Strategy

After the QP is solved, the spatial deviation between each smoothed trajectory point and its reference waypoint is computed. If the fraction of points exceeding the safe corridor boundary surpasses 20%, the Bézier result is discarded, the original reference path is retained, and kinematic quantities are re-estimated via finite differences. QP convergence failures arise when the RRT* detour path deviates far enough from the beam axis that the corridor bounds and the deviation penalty impose competing demands, shrinking the feasible domain. If the QP itself fails to converge, the solution transitions to the linear interpolation initial guess and kinematics are estimated through the same procedure.

The guarantee provided by this mechanism is bounded in scope. When a feasible QP solution is found, the sampled trajectory points satisfy the rectangular corridor inequality at the discrete 10 s time steps; continuous-time containment within the circular beam cross-section is not guaranteed, for the reasons discussed in Section 4.4.3. When the QP is infeasible or the exceedance ratio surpasses 20%, the guarantee is limited to trajectory availability: an executable path is always produced, but beam coverage may be compromised. The coverage rates reported in Sections 5.2–5.5 are computed at discrete sampling instants and should be interpreted within this scope.

5. Simulation Results and Analysis

5.1. Experimental Setup

5.1.1. Simulation Platform and Environment

The simulation experiments were conducted on MATLAB R2024a. Satellite orbit propagation was performed using the SGP4 model based on standard two-line element (TLE) data for time-series prediction and visible arc computation. All experiments were executed under the same hardware and software environment to ensure reproducibility.

5.1.2. System Parameter Configuration

The Wuhan ground station (30.5279° N, 114.888° E) was adopted as the observation reference. A Starlink satellite (NORAD ID 44717, orbital inclination 53°, orbital altitude approximately 550 km, with an orbital period of approximately 96 min) was selected as the target satellite. The simulation time window spanned from 00:00 to 12:00 UTC on 14 January 2025, with a time step of 10 s. The UAV was assigned a fixed flight altitude of 1500 m, a maximum speed of 50 m/s, and a maximum acceleration of 5 m/s². The minimum elevation threshold of the radar was set to 30°, and the beam half-cone angle was 2°. The main simulation parameters are summarized in Table 2. The no-fly zone configuration is listed in Table 3.

The RRT* parameters in Table 2 balance exploration breadth against search speed. The goal bias $p_{\text{goal}} = 0.12$ and geometry bias $p_{\text{geom}} = 0.25$ direct 37% of samples toward the goal and the beam axis, while the remainder explore uniformly, a split that performed reliably across the obstacle densities in Sections 5.3–5.5. The 20 m extension step equals roughly 20% of the minimum corridor radius at $\epsilon = 30^\circ$, keeping individual extensions

small relative to the available detour space. The Bézier weights $\lambda_{\text{snap}} = 1.0$ and $\lambda_{\text{dev}} = 10.0$ prioritize beam containment over smoothness by penalizing reference-path deviation more heavily than snap energy; the sensitivity of this choice is examined in Section 5.6.2. The 20% fallback threshold provides a conservative margin.

5.1.3. Performance Metrics

The performance metrics adopted in this study follow standard evaluation practices in UAV coverage path planning [9]. The following metrics are adopted for evaluation of the path planning and trajectory optimization results:

1. Beam coverage rate: The fraction of time steps at which the spherical angular distance between the UAV and satellite line-of-sight vectors falls within the beam half-cone angle of 2° , expressed as a percentage.
2. Azimuth and elevation tracking error: The absolute difference in azimuth and elevation angles between the UAV and the satellite at each time step, by which the angular alignment accuracy is quantified.
3. Maximum speed and acceleration: The peak instantaneous speed and acceleration along the trajectory, verified against the prescribed kinematic limits of 50 m/s and 5 m/s^2 .
4. Obstacle avoidance success rate: The proportion of trajectory points that remain outside all no-fly zone boundaries, with 100% indicating complete avoidance.

5.2. Obstacle-Free Scenario: Results and Analysis

The obstacle-free scenario serves as the baseline validation case for the proposed method. In this scenario, no no-fly zones are present in the flight airspace, and three constraints are imposed: a kinematic constraint limiting the speed to 50 m/s and the acceleration to 5 m/s^2 , a beam geometric constraint requiring G-U-S collinearity, and a satellite visibility constraint requiring the elevation angle to remain at or above 30° . This scenario corresponds to the system schematic shown in Figure 2. A 12-h TLE data window was scanned to identify all visible passes of the target satellite above the 30° elevation threshold. Multiple visible arcs were detected, with individual durations ranging from approximately 2 to 5 min depending on the satellite's elevation profile during each pass. The segment from 04:07 to 04:10 (approximately 3.3 min) is reported here as a representative example. The same four-stage pipeline was applied identically to every detected arc; coverage and kinematic profiles across the remaining arcs were consistent with the representative case, and no arc triggered the fallback mechanism under the nominal three-obstacle configuration. Throughout this visible arc, a fixed altitude is maintained by the UAV, and the beam axis directed from the ground station toward the satellite is followed, thereby satisfying the collinear condition. This scenario is used to verify whether complete beam coverage can be achieved under ideal conditions.

To evaluate the azimuth and elevation tracking performance, the angular profiles of the satellite and the UAV over the representative visible arc are plotted in Figure 5. The satellite elevation followed a characteristic ascending-then-descending arc, peaking at approximately 75° around 04:08:30 and remaining above the 30° minimum observable elevation throughout. The satellite azimuth underwent a coordinate wrap-around after zenith passage, crossing from approximately 350° to approximately 40° as a normal quadrant transition in angular coordinates. Close overlap was observed between the UAV azimuth and elevation curves and those of the satellite throughout the visible arc, confirming that the G-U-S collinear constraint was effectively maintained, and satisfying the tracking accuracy requirement defined in Section 5.1.3.




Figure 5. Azimuth and elevation profiles of the satellite and the UAV over a visible arc. The UAV curves closely overlap those of the satellite, verifying the G-U-S collinear constraint.

Figure 6 presents the instantaneous speed and acceleration profiles over the visible arc. The speed followed a descending-then-ascending pattern: it started at approximately 28 m/s at the beginning of the pass, decreased gradually to a minimum of approximately 22.5 m/s near the zenith around 04:08:30, and then recovered to approximately 27.5 m/s as the satellite elevation declined. This pattern is consistent with the variation in the satellite angular velocity, since a higher elevation produces a smaller angular velocity projected onto the ground, reducing the tracking speed required of the UAV. The peak speed over the entire arc was approximately 28 m/s, well below the 50 m/s constraint. The acceleration increased monotonically from approximately -0.12 m/s^2 to approximately $+0.11 \text{ m/s}^2$, with a maximum absolute value of approximately 0.12 m/s^2 , far below the 5 m/s^2 limit. The low acceleration level indicates that the satellite angular velocity varied gently in this scenario, and only minor maneuvering was required of the UAV to maintain alignment; the remaining acceleration margin is available for handling sudden disturbances or more complex orbital geometries. Both metrics were verified against the kinematic feasibility requirement defined in Section 5.1.3.




Figure 6. UAV kinematic profiles over the visible arc: (a) instantaneous speed; (b) instantaneous acceleration. The peak speed and acceleration are approximately 28 m/s and 0.12 m/s^2 , both well within the prescribed limits.

Complete beam coverage was achieved at every time step over the visible arc, as confirmed in Figure 7. At each time step, the spherical angular distance between the UAV line-of-sight vector and the satellite line-of-sight vector was computed; the UAV was considered within the beam when this distance did not exceed the beam half-cone angle of 2° . A coverage rate of 100.0% was obtained, consistent with expectation, since the UAV operates along the beam axis throughout the pass and no factor forces the trajectory to deviate from the safe corridor.




Figure 7. Beam coverage status over the visible arc. The UAV remains within the beam at all time steps, yielding a coverage rate of 100.0%.

Figure 8 compares the planned trajectory with the ideal reference path in geographic coordinates. The ground station is marked by a red pentagram at approximately 30.5279° N, 114.888° E. The ideal trajectory is shown as a blue solid line and the planned trajectory as a green dashed line; the two curves nearly overlap throughout the pass. The trajectory extended from southwest to northeast over a spatial extent of approximately 4 km, consistent with the satellite transit direction. All trajectory points were marked by green filled circles, indicating beam coverage at every sampled instant and corroborating the 100.0% coverage rate. The close agreement demonstrates that the Minimum Snap Bézier optimization successfully smoothed the discrete ideal waypoints into a continuous flyable trajectory under the safe corridor constraint, with the position adjustments introduced by smoothing remaining well within the allowable beam radius.



Figure 8. Trajectory comparison in geographic coordinates for the obstacle-free scenario. Green dots indicate UAV positions that fall within the radar beam; the red star marks the ground station. The ideal trajectory and the planned trajectory nearly overlap throughout the pass, with all points remaining within beam coverage.

These results establish the performance ceiling of the proposed framework. The 100% coverage rate demonstrates that the ideal position computation and Minimum Snap smoothing preserve beam alignment at every time step when obstacles are absent. The peak speed reaches only 56% of the 50 m/s limit and the peak acceleration remains below 3% of the 5 m/s^2 limit, indicating ample maneuverability reserve for disturbances or more demanding orbital geometries. This baseline also validates the ideal position generation method (Section 3.3.1) independently of the RRT* module, isolating the framework's two core functions.

5.3. Obstacle Scenario: Results and Analysis

To evaluate the effectiveness of the geometry-guided RRT* algorithm in complex airspace, three circular no-fly zones were placed along the ideal trajectory within the first visible arc. Their radii are 45 m, 55 m, and 40 m, with centers located at 28%, 52%, and 74% of the ideal trajectory arc length, respectively. All lateral offsets were set to zero, so that each center lies directly on the ideal path, creating a direct obstruction to the collinear trajectory. The time step was set to 10 s. When the ideal trajectory intersected a no-fly zone, the RRT* obstacle-avoidance planning mode was automatically activated, and a feasible detour path was searched under the beam geometry guidance; the resulting discrete path was then smoothed into an executable trajectory by the Minimum Snap Bézier optimization. In practical monitoring missions, 100% continuous beam coverage throughout the entire visible arc is not required; a cumulative coverage rate exceeding the mission-specific threshold during the critical observation period is sufficient to meet the demands of data acquisition, relay communication, or target tracking.

Figure 9 shows the azimuth and elevation profiles under the obstacle scenario. The UAV curves remained closely aligned with those of the satellite throughout the pass. The elevation rose continuously from approximately 30° to approximately 70° , reflecting the typical geometry of a satellite transit. As in the obstacle-free case, the satellite azimuth underwent a normal coordinate wrap-around near the zenith, crossing from approximately 350° to approximately 40° . Despite the local detours introduced in the obstacle mode, the azimuth and elevation curves retained high consistency with the obstacle-free scenario, because the horizontal deviation caused by obstacle circumvention is negligible relative to the satellite-to-ground distance. This result confirms that high-precision beam alignment is maintained even while no-fly zones are being actively avoided, satisfying the tracking accuracy requirement defined in Section 5.1.3.




Figure 9. Azimuth and elevation profiles of the satellite and the UAV over the visible arc under the obstacle scenario. The UAV curves remain closely aligned with those of the satellite despite local detours, verifying that the beam alignment is maintained during obstacle avoidance.

Figure 10 presents the speed and acceleration profiles under the obstacle scenario. The speed varied within 18–26 m/s, corresponding to only 36–52% of the 50 m/s constraint, while the acceleration fluctuated between -0.7 and 0.5 m/s^2 , remaining below 15% of the 5 m/s^2 limit. The speed profile exhibited local fluctuations during the detour segments. A trough of approximately 18 m/s around 04:09:20 corresponded to the obstacle-avoidance maneuvers as the three no-fly zones were sequentially bypassed: a moderate deceleration was applied when approaching each zone to adjust the heading, followed by re-acceleration under the G-U-S geometric guidance to rejoin the ideal trajectory after clearance. The acceleration profile showed a corresponding negative-then-positive oscillation during the same intervals, consistent with the speed variation. The Minimum Snap Bézier optimization effectively eliminated the velocity discontinuities introduced by the discrete RRT* waypoints, confining the dynamic response during the detour transitions within a narrow range. Both metrics were verified against the kinematic feasibility requirement defined in Section 5.1.3.




Figure 10. UAV kinematic profiles over the visible arc under the obstacle scenario: (a) instantaneous acceleration; (b) instantaneous speed. The speed varies within 18–26 m/s and the acceleration fluctuates between -0.7 and 0.5 m/s^2 , both well within the prescribed limits.

Figure 11 plots the UAV angular positions against the radar beam coverage region in the azimuth-elevation space. The elliptical outlines represent the radar beam coverage region at each sampling instant for a half-cone angle of 2° . A coverage rate of 95.2% was achieved over the full arc. During the non-detour phases, the UAV angular positions fell precisely within the beam boundary, confirming that the positioning accuracy satisfied the 2° half-cone angle requirement. After each obstacle-avoidance maneuver, the system rapidly recovered an effective coverage state. This rapid return of the trajectory within the beam coverage region after circumvention validates the critical role of the constrained embedding in enforcing convergence within the corridor boundary, consistent with the expectation established in Section 5.2.

The spatial relationship between the planned trajectory and the no-fly zones is presented in Figure 12. The blue curve denotes the obstacle-free ideal trajectory, the green dashed curve represents the RRT* obstacle-avoidance path, the red circular regions mark the no-fly zones, and the green and red scatter points indicate time steps within and outside the beam coverage, respectively. The RRT* planned path completely circumvented all no-fly zones, with no trajectory point intruding into any restricted boundary, yielding a 100% obstacle avoidance rate. After each detour maneuver was completed, the path converged rapidly toward the ideal trajectory under the G-U-S geometric guidance. The overall flight region was confined within several kilometers around the ground station, and the trajectory remained continuous and smooth throughout.




Figure 11. Beam coverage in the azimuth-elevation space under the obstacle scenario. The UAV remains within the beam during non-detour phases and recovers coverage rapidly after each avoidance maneuver, yielding a full-arc coverage rate of 95.2%.



Figure 12. Trajectory comparison in geographic coordinates under the obstacle scenario. Red circles, crosses, and the star denote no-fly zones, their centres, and the ground station, respectively. The RRT* path circumvents all obstacles and converges rapidly toward the ideal trajectory, achieving 95.2% beam coverage.

Taken together, these results demonstrate that the geometry-guided RRT* resolves the central tension identified in Section 1: it achieves 100% obstacle avoidance while holding the coverage loss to 4.8 percentage points relative to the obstacle-free baseline. The rapid reconvergence toward the ideal path after each detour confirms that the G-U-S geometric guidance remains effective even when the beam axis is briefly departed from for obstacle circumvention. This coverage loss reflects a physical limit: a detour around an obstacle placed on the beam axis cannot also occupy the axis. Whether such a loss is acceptable depends on the mission-specific coverage threshold discussed in Section 5.1.3. The multi-run characterization in Section 5.6.4 provides statistical context for the single-run values reported here.

5.4. Comparative Analysis

To assess the contribution of the geometry-guided sampling strategy to the overall framework, an ablation experiment was conducted in which the geometry-guided RRT* was replaced with standard RRT* under uniform random sampling. The goal bias and corridor bias were both set to zero, so that sampling was spatially uniform with no preference for regions near the goal or the beam axis. All other components of the framework, including the Minimum Snap Bézier QP optimization with the conservative corridor constraint, remained unchanged.

Under this ablated configuration, standard RRT* achieved a beam coverage rate of 85.7%, compared with 95.2% for the proposed geometry-guided RRT*. The 9.5-percentage-point gap is attributable to the uniformly sampled paths deviating sufficiently from the ideal trajectory to cause the Bézier QP optimization to exceed the 20% out-of-beam threshold and trigger the fallback mechanism. Once triggered, the fallback reverts to finite-difference kinematics computed directly from the raw RRT* path, forfeiting the smoothness and kinematic feasibility guarantees that the QP optimization would otherwise provide. The geometry-guided RRT* run discussed here remained within the 20% deviation threshold and did not trigger the fallback. Section 5.6.4 reports multi-run data showing that both fallback frequency and coverage vary with the random seed, as expected for stochastic RRT* sampling.

This result confirms that the geometry-guided sampling strategy is not an incremental refinement but an essential component of the framework. Without goal-biased and corridor-constrained sampling, the RRT* planner produces waypoints that lie outside the convergence basin of the subsequent QP optimization, forcing the system to rely on a degraded fallback trajectory. The sampling strategy and the hard-constrained optimization are therefore co-designed: the former supplies waypoints that the latter can feasibly smooth within the beam corridor, and neither component alone is sufficient to deliver both beam containment and kinematic feasibility.

To contextualize the proposed framework within the broader planning landscape, Table 4 provides a qualitative comparison with three representative alternative planner classes: Informed RRT*, Dubins-based fixed-wing planning, and A* with quadratic programming post-processing. Each method is evaluated on five dimensions critical to the satellite monitoring mission. The proposed geometry-guided RRT* is the only approach that natively addresses all five: it incorporates beam alignment as a planning prior (via the G-U-S ideal positions), handles obstacles through sampling-based detours, respects fixed-wing kinematics through the Bézier QP stage, embeds a conservative beam corridor constraint, and includes a diagnostic fallback mechanism that guarantees trajectory availability under all conditions. Informed RRT* shares the RRT* foundation but its ellipsoidal sampling heuristic is driven by path length minimization rather than beam-axis adherence, making it less effective when the corridor centerline is occluded (as demonstrated in the boundary edge case, Section 5.5.1). Dubins-based planners natively satisfy the minimum-turn-radius constraint of fixed-wing aircraft but are designed for 2D path-length optimization in static environments and do not incorporate a time-varying beam corridor. A* with QP post-processing provides deterministic, resolution-optimal paths and can enforce corridor constraints at the QP stage, but the graph-search front-end lacks the geometric guidance to keep waypoints within the corridor convergence basin, making it susceptible to the same QP feasibility failures observed in the uniform-sampling RRT* ablation.

A quantitative coverage comparison across all five planner classes is not reported because a fair evaluation requires each alternative method to be implemented, tuned, and tested within the identical G-U-S simulation environment. Such a multi-planner benchmark

constitutes a substantial standalone study and is deferred to future work. The ablation experiment reported in this section isolates the specific contribution of the geometry-guided sampling by removing it while holding all other pipeline components constant. This controlled comparison directly answers the question of whether the guidance strategy matters: the 9.5-percentage-point coverage gap between uniform and geometry-guided RRT* affirmatively resolves it. The qualitative comparison in Table 4 further contextualizes the proposed framework within the broader planning landscape, highlighting the unique combination of beam alignment, obstacle avoidance, kinematic handling, and diagnostic fallback that no single alternative method currently provides.

Table 4. Qualitative comparison of planning approaches for UAV-assisted LEO satellite monitoring.

Capability	Proposed	Informed RRT*	Dubins	A* + QP	Uniform RRT*
Beam alignment as planning prior	✓	–	–	–	–
Obstacle avoidance	✓	✓	✓	✓	✓
Fixed-wing kinematic constraints	✓	–	✓	✓ [†]	–
Conservative beam corridor (QP)	✓	✓ [*]	–	✓ [*]	✓ [*]
Diagnostic fallback	✓	–	–	–	–
Deterministic output	–	–	✓	✓	–

✓ = natively supported; – = not supported. ^{*} Applicable only if a QP post-processing stage is appended to the planner. [†] Kinematic feasibility enforced at the QP stage only; the front-end path may violate turn-radius constraints.

5.5. Edge Case Stress Testing

Two edge case scenarios, beyond the three-obstacle baseline of Section 5.3, probe the framework’s behavior under extreme conditions.

5.5.1. Boundary Obstacle at Minimum Corridor Width

The framework is most constrained at the start of each visible arc. Satellite elevation hovers near the 30° threshold, and the safe corridor radius drops to its minimum: $R_c \approx 104.7$ m at $H = 1500$ m, $\varepsilon = 30^\circ$. A single circular no-fly zone of radius 55 m was placed at relative arc-length position 0.04, directly on the ideal trajectory, to stress this limiting condition.

The RRT* planner could not discover a complete collision-free detour within the 2000-iteration budget. At this elevation, the corridor offers a total lateral clearance of approximately 209 m (diameter) against an obstacle diameter of 110 m, leaving a usable margin of roughly 99 m. The margin exceeds the 20 m RRT* extension step, so the geometry is not the bottleneck. Rather, the geometry guidance bias directs candidate nodes toward the beam axis, which the obstacle itself occludes, and the uniform-sampling component lacks the density needed to discover the narrow viable passage within the iteration budget. The fallback triggered: Bézier QP optimization was skipped, the reference path was retained with finite-difference kinematics, and the resulting trajectory delivered 100.0% beam coverage with 95.2% obstacle avoidance.

We note that this result does not reflect geometric impossibility. The corridor diameter of 209 m exceeds the obstacle diameter of 110 m by 99 m, so a lateral passage exists. What fails is the co-design: the sampling bias that accelerates planning under nominal conditions becomes counterproductive when the beam axis is occluded, and the 2000-iteration budget is insufficient for the uniform component to compensate. The fallback preserves trajectory availability, but the mission operator must accept a trade-off: either a brief coverage gap if the QP is retained or a minor obstacle incursion if the fallback path is used.

5.5.2. Combined Boundary and Mid-Arc Obstacles

Adding the boundary obstacle (relative position 0.04, radius 55 m) to the original three mid-arc obstacles (positions 0.28, 0.52, 0.74) produced a worst-case configuration that simultaneously stresses the narrowest corridor point and demands multiple downstream detours. While the baseline experiments (Sections 5.2–5.4) used a single representative arc, the edge cases were run across the full 12-h TLE scan window. The RRT* planner triggered detours for seven of the obstructed segments; the Bézier QP fallback activated. Coverage reached 76.2% and obstacle avoidance 85.7%, the lower bound of observed performance across all tested configurations. The 19.0 percentage-point drop from the three-obstacle baseline (95.2%) is driven by the boundary obstacle, which forces a lateral deviation at the narrowest corridor segment where the feasible detour region is smallest.

Table 5 collects coverage and obstacle avoidance rates for all configurations into a single performance envelope.

Table 5. Performance envelope of the proposed framework across obstacle configurations.

Configuration	NFZs	Coverage (%)	Obs. Avoid. (%)
Obstacle-free (baseline)	0	100.0	100.0
Original 3-obstacle (baseline)	3	95.2	100.0
Boundary single obstacle	1	100.0	95.2
Boundary + 3 mid-arc obstacles	4	76.2	85.7

Taken together, these two edge cases clarify the framework’s primary vulnerability. It is not geometric: a feasible detour exists around every obstacle tested, and the corridor diameter exceeds the largest obstacle by nearly a factor of two. The vulnerability is algorithmic. Geometry guidance concentrates sampling near the beam axis; when that axis is occluded, the sampling budget that produced efficient plans under nominal conditions becomes insufficient to explore lateral passages. This tension is inherent to any guided-sampling strategy and cannot be eliminated by parameter tuning alone. The fallback mechanism (Section 4.4.4) guarantees that a flyable trajectory is always produced, but the trade-off between coverage and obstacle avoidance that emerges under extreme occlusion (100.0% coverage with 95.2% avoidance for the boundary obstacle; 76.2% coverage with 85.7% avoidance for the combined case) must be accounted for in mission planning. The failure mode catalog in Section 5.6.4 generalizes this analysis to four additional operating dimensions: high satellite angular velocity, extremely short visible arcs, GPS-denied environments, and large-scale obstacle fields.

5.6. Discussion

5.6.1. Effect of Time Step on Constraint Satisfaction

To evaluate the influence of the time step on path planning performance, simulations were conducted under the obstacle scenario at four time step settings: 30 s, 20 s, 15 s, and 10 s. The results are shown in Figure 13. The coverage rate increased monotonically as the time step was reduced, rising from 57.1% at 30 s to 90.9% at 20 s, 92.9% at 15 s, and 95.2%

at 10 s. The improvement gradually leveled off at the 10 s step, indicating that this value is close to the coverage upper bound for the given scenario.




Figure 13. Effect of time step on beam coverage rate under the obstacle scenario: (a) 30 s, coverage 57.1%; (b) 20 s, coverage 90.9%; (c) 15 s, coverage 92.9%; (d) 10 s, coverage 95.2%.

This trend arises from three coupled mechanisms. The first is a detection density effect, in which a larger time step reduces the total number of coverage checkpoints within the visible arc, amplifying the contribution of the few uncovered points in detour segments to the overall ratio and thereby lowering the reported coverage rate. The second is a kinematic constraint relaxation effect, in which a larger time step equivalently permits greater single-step displacement under the speed constraint, granting the RRT* detour path more room to deviate but simultaneously weakening the kinematic feasibility of returning within the beam promptly after circumvention. The third is a control point sparsity effect, in which decreasing the number of B'ezier control points as the time step increases reduces the precision with which the trajectory adheres to the corridor boundary. Under the simulation conditions adopted in this study, the highest coverage rate is achieved at a time step of 10 s. In practical deployment, the time step should be jointly determined by the visible arc duration and the maximum UAV speed to ensure that full beam alignment is completed by the trajectory between consecutive control points.

5.6.2. Effect of Deviation Penalty Weight on Optimization Stability

Whereas the time step governs the temporal resolution of the planning pipeline, the optimization behavior is also sensitive to the internal weighting of the objective function. The objective function of the Minimum Snap B'ezier optimization consists of a snap minimization term and a corridor deviation penalty term, with λ_{dev} serving as the deviation penalty weight. To examine its influence on optimization stability, comparative experiments were conducted under the obstacle scenario with the time step fixed at 10 s, using $\lambda_{\text{dev}} = 10$

and $\lambda_{\text{dev}} = 50$. At $\lambda_{\text{dev}} = 10$, a stable solution was reached by the optimization and a coverage rate of 95.2% was attained. At $\lambda_{\text{dev}} = 50$, the coverage rate exhibited pronounced instability: three independent runs under identical conditions yielded coverage rates of 66.7%, 83.3%, and 95.2%, respectively (Figure 14). A deterministic solution is therefore not reached when the deviation penalty weight is excessively large.




Figure 14. Beam coverage results under the obstacle scenario with $\lambda_{\text{dev}} = 50$ (time step = 10 s): (a) Run 1, coverage 66.7%; (b) Run 2, coverage 83.3%; (c) Run 3, coverage 95.2%. The three runs yield substantially different coverage rates, illustrating the instability of the soft-constraint penalty approach at high λ_{dev} .

This instability originates from the optimization entering a boundary-competition regime. In the presence of no-fly zones, the trajectory is subject to two competing forces: an obstacle repulsion, generated by the feasible domain boundaries imposed through collision constraints, and a corridor attraction, driven by the deviation penalty term weighted by λ_{dev} . When λ_{dev} is small, the deviation penalty exerts limited influence relative to the snap term, the optimizer is dominated by the smoothness objective, and a unique stable solution is obtained within the feasible domain. As λ_{dev} is increased, the corridor attraction and the obstacle repulsion become comparable in magnitude, causing the gradient direction to reverse repeatedly near the feasible domain boundary. The active constraint set of the QP thereby alternates rapidly between successive iterations, preventing the solver from converging to a deterministic optimum and producing the observed fluctuations in coverage rate.

This sensitivity to λ_{dev} validates the hard-constraint embedding strategy adopted in the proposed framework: by formulating the beam coverage condition directly as linear inequality constraints rather than relying on a penalty term alone, the boundary-competition problem is eliminated at its root. Together with the ablation experiment, which demonstrates that geometry-guided sampling is essential for supplying waypoints within the feasible domain, this analysis confirms that the sampling strategy and the constrained optimization form a co-designed system in which each component addresses a distinct failure mode.

5.6.3. Toward Practical Deployment

The present study is simulation-based, but the pipeline can be deployed on a physical fixed-wing UAV with the architecture shown in Figure 15. Five subsystems cooperate: an onboard computer executes the four-stage planner before takeoff; an autopilot handles real-time state estimation, trajectory tracking, and actuator control; the LEO satellite ephemeris is propagated from publicly available TLE data; a ground station provides mission configuration and monitoring; and environmental disturbances are absorbed by the beam corridor margin. The ephemeris computation and ground station setup follow Sections 3.3.1 and 5.1. We examine three engineering dimensions: onboard computation, autopilot integration,

and disturbance accommodation. We then outline a progressive validation sequence that would transition the framework from simulation to flight testing.




Figure 15. UAV onboard architecture and communication loops for practical deployment. (a) Pre-mission computation pipeline: TLE propagation, RRT* path planning, and Bézier QP optimization executed on the onboard computer before takeoff; (b) autopilot communication loops: MAVLink-based waypoint upload and execution via PX4/ArduPilot; (c) LEO target satellite with TLE-based ephemeris; (d) ground station for mission parameter configuration and monitoring; (e) environmental disturbances (GPS error, wind) accommodated by the beam corridor margin.

Onboard computational architecture. The planning pipeline is structured to run on an embedded computer mounted on the UAV (Figure 15a). SGP4 propagation (Stage 1) and coverage assessment (Stage 4) are lightweight: each requires only analytical function evaluations and contributes negligibly to the overall budget. The RRT* path planner (Stage 2) dominates the workload. Its sampling-based search iterates per segment with repeated collision checks and nearest-neighbor queries, and we measured 1.35 s across six obstructed segments (0.22 s per segment at 2000 iterations) on our desktop platform. The Bézier QP (Stage 3) solves in 2.57 s for $n_c \approx 360$ decision variables, where the dimension scales as $n_c = 3 \times (n_b + 1) \times (N_{wp} - 1)$. The Hessian matrix dominates memory, occupying roughly 1 MB of working memory, well within the capacity of current embedded processors, for which interior-point and active-set solvers are commercially available.

Measured on a desktop machine (MATLAB R2024a, nominal three-obstacle configuration), the core pipeline completes in approximately 5.6 s per visible arc: 1.7 s for SGP4 propagation over the 12-h scan window, 1.35 s for RRT* planning, 2.57 s for the QP, and under 0.1 s for coverage assessment. We emphasize that these figures represent a pre-mission computation, not an in-flight one. Because the offline paradigm (Section 4.1) generates the complete trajectory before takeoff, only setpoint streaming to the autopilot is required during flight. The 10 s planning time step leaves a margin of roughly 4.4 s, which provides comfortable headroom for the onboard processor even after accounting for the embedded-to-desktop performance gap. In-flight replanning, if triggered by updated TLE data or a newly detected obstacle, would re-execute the same pipeline on the onboard computer.

Autopilot integration and communication. The framework outputs time-stamped ENU positions and velocities that map directly to the waypoint-command interface of PX4, ArduPilot, and similar open-source autopilots (Figure 15b). Position commands travel as SET_POSITION_TARGET_LOCAL_NED MAVLink messages; velocity feed-forward terms from the Bézier evaluation accompany them as TWIST messages, which measurably improve tracking on curved segments. A thin ROS 2-to-MAVLink bridge translates planning output

to flight-stack commands. Two data streams operate in parallel: a low-rate uplink carries ephemeris updates, no-fly zone revisions, and diagnostic commands; a downlink reports telemetry and coverage status. Because the complete trajectory is pre-loaded, a loss of the uplink does not interrupt the mission; the UAV simply continues executing the stored plan or, if necessary, transitions to a predefined fail-safe mode.

Environmental disturbance accommodation. The beam safety corridor (Equation (6)) supplies an inherent geometric margin. At the most constraining elevation, $\varepsilon = 30^\circ$, the corridor radius reaches $R_c \approx 104.7$ m. Four uncertainty sources compete for this margin: GPS positioning error (roughly 5 m CEP for civilian receivers), wind-induced lateral drift (10–20 m between consecutive control updates at 1500 m altitude), autopilot tracking error, and communication latency. The corridor radius exceeds the combined GPS and wind contribution by a factor of approximately four, leaving positive headroom for the remaining terms under moderate conditions. The 20% fallback threshold (Algorithms 1 and 2) provides a second safety layer. For persistently high-wind environments, a narrower half-cone angle substitutes directly into Equation (6). We regard active disturbance rejection, implemented as a model predictive control wrapper that re-triggers planning upon detecting a sustained deviation, as a priority for future flight testing.

Progressive validation roadmap. We envision a four-step transition from the present simulation study to flight-proven operation, each step isolating a distinct integration risk. Software-in-the-loop validation would first confirm MAVLink message-level compatibility by feeding the pipeline output into a desktop autopilot simulator running PX4 or ArduPilot firmware. Hardware-in-the-loop testing would then deploy the pipeline onto the target embedded computer, connected to a physical autopilot over a serial link, to measure data latency and trajectory tracking fidelity under benchtop conditions. A constrained outdoor flight campaign would follow, exercising the geometry-guided RRT* obstacle avoidance and Bézier trajectory tracking against static no-fly zones under real flight dynamics. This is the first point at which sensor noise, wind, and autopilot transients enter the loop. Finally, a live satellite pass with real-time TLE input would exercise the complete G-U-S alignment pipeline end to end, providing a direct measurement of beam coverage performance on physical hardware. These four steps verify planner correctness, communication reliability, flight control accuracy, and beam alignment in order of increasing integration complexity.

5.6.4. Limitations and Failure Modes

The boundary obstacle experiment in Section 5.5 establishes one failure mode directly: corridor saturation. Four further conditions, derived from the kinematic and environmental constraints in the system model (Section 3.3.2), complete the catalog below. We present each mode with its triggering condition, observed or projected impact, and a candidate mitigation.

Corridor saturation. This is the only failure mode we have observed experimentally. When an obstacle occupies the beam axis at the narrowest corridor point, the geometry-biased sampling directs candidate nodes toward the blocked region; the uniform-sampling component, at 15% probability in the current configuration, cannot reliably discover the lateral passage within the 2000-iteration budget, even though the passage exists geometrically (Section 5.5). The fallback triggers, forcing a coverage-avoidance trade-off. Adaptive step-size scaling that widens the sampling region as the corridor narrows is a natural countermeasure.

Extremely short visible arcs. A satellite pass lasting under one minute produces fewer than three waypoints at the 10 s time step. The Bézier order then defaults to cubic, but with so few segments the trajectory cannot resolve the beam axis curvature. Multi-resolution

waypoint placement, which inserts synthetic intermediate anchors below the 10 s sampling floor, would address this regime.

High satellite angular velocity. Orbits below roughly 300 km project angular velocities onto the UAV altitude plane that exceed the 50 m/s platform limit. Here the kinematic constraint (Equation (4)), not the planning algorithm, is the bottleneck. Variable-altitude planning, which adjusts the UAV altitude to trade corridor radius for required speed, offers a partial remedy.

GPS-denied or TLE-unavailable environments. The G-U-S collinear prior collapses when either the satellite ephemeris or the ground station position is unavailable. The pipeline cannot start. An onboard angle-of-arrival estimator that locks onto the satellite beacon directly, bypassing the SGP4 propagation chain, would restore operation in GPS-compromised settings.

Large-scale obstacle fields. The geometry-guided RRT* targets the sparse obstacle densities typical of low-altitude airspace. Urban canyons and building-scale no-fly zones present a qualitatively different environment: corridor-constrained samples are predominantly blocked, and the 15% uniform-sampling budget becomes too thin to explore viable routes. Raising the uniform-sampling fraction or switching the front-end to a graph-search planner (e.g., A*) within the corridor are candidate adaptations.

These five conditions mark the edges of the framework's current applicability. Each is accompanied by a concrete extension path, and none represents a fundamental defect: the geometric-guidance principle remains sound; the boundaries arise from the specific sampling and kinematic parameter choices made for the nominal operating regime.

Two broader considerations further delineate the scope. First, the current formulation handles a single UAV tracking a single satellite. Multi-UAV cooperative monitoring would require a task assignment layer that allocates visible arcs across the fleet, since the G-U-S collinear condition is defined per triplet. Second, the RRT* computational cost scales with obstacle count rather than environment size, because sampling remains confined to the beam corridor neighborhood. For the sparse densities anticipated in low-altitude airspace, this cost is manageable without algorithmic revision.

Stochastic variability. Twenty independent trials per configuration, each with a distinct random seed, were run under the nominal parameter settings. The three-obstacle baseline produced a mean coverage of $86.2 \pm 9.3\%$ (range 66.7–100.0%); the combined boundary-plus-three configuration produced $84.0 \pm 7.8\%$ (range 61.9–95.2%). The single-run values reported in Sections 5.2–5.5 (95.2% and 76.2%) fall within these ranges and represent different points in their respective distributions. With an 85% coverage threshold for direct QP success, the baseline cleared the threshold in 14 of 20 runs and the combined case in 11 of 20. In every remaining run, the fallback triggered: an executable trajectory with reduced coverage but full obstacle avoidance was produced. No run failed to yield a flyable trajectory.

We emphasize one implication of these data. The 33.3-percentage-point spread in the baseline is not a QP instability. It reflects the architectural role of the fallback: when a random seed places RRT* waypoints away from the beam axis, the QP feasible domain shrinks and the fallback activates, producing a lower-coverage trajectory that is nevertheless executable. The spread is driven by stochastic sampling, not by optimization behavior, a finding consistent with the boundary-competition analysis in Section 5.6.2. The coverage rates in Sections 5.2–5.5 should therefore be read as indicative performance levels, not deterministic guarantees. A 30-seed Monte Carlo study would supply the confidence intervals needed for mission-grade reliability assessment.

6. Conclusions

This paper proposed a UAV-assisted path planning method that integrates beam geometry guidance with constrained trajectory optimization. Centered on the G-U-S collinear geometric constraint, the framework unifies visible arc extraction, dual-mode path planning, and Minimum Snap optimization under a conservative corridor constraint. The work presents a simulation framework for the brief visibility window that terrain occlusion and Earth curvature impose on ground-based LEO satellite monitoring.

Together, the simulation results establish a consistent pattern. Obstacle-free flight delivers 100% beam coverage at speeds and accelerations well within platform limits, validating the ideal-position computation independently of the RRT* module. With three no-fly zones placed on the beam axis, geometry-guided RRT* achieves 95.2% coverage and complete obstacle avoidance. Removing the geometry guidance drops coverage by 9.5 percentage points and triggers the diagnostic fallback; the sampling strategy is therefore load-bearing, not incremental. The time-step experiment shows that coverage degrades monotonically as the step widens, confirming that the 10 s sampling density is necessary for the reported performance. The λ_{dev} experiment reveals that a large soft-penalty weight drives the QP into a boundary-competition regime where the solution oscillates across runs; the hard-inequality formulation eliminates this instability at its source.

Edge case testing identifies the framework's primary vulnerability. A boundary obstacle placed at the minimum corridor width ($\varepsilon \approx 30^\circ$) triggers the fallback, exposing a tension in the co-design: the geometry bias that accelerates nominal planning becomes an exploration bottleneck when the beam axis itself is occluded. The fallback preserves trajectory availability, but the coverage-avoidance trade-off that emerges under extreme occlusion must be factored into mission planning: 100.0% coverage with 95.2% avoidance for the boundary case, and 76.2% coverage with 85.7% avoidance for the combined case. Five failure modes, catalogued in Section 5.6, each carry a concrete extension path. None indicates a fundamental defect in the geometric-guidance principle; the boundaries arise from parameter choices made for the nominal operating regime. Taken together, the simulation results support a bounded conclusion: under the configurations tested, the framework produces executable trajectories that simultaneously satisfy beam alignment, obstacle avoidance, and kinematic limits, with the fallback mechanism ensuring trajectory availability regardless of sampling outcome.

Four conditions bound the present results: fixed flight altitude, a single satellite, one representative visible arc reported in detail, and a desktop simulation platform. Extending to three-dimensional planning, multi-satellite cooperative task assignment, and hardware-in-the-loop validation are the natural next steps. The failure mode analysis points to adaptive step-size scaling for narrow corridors and sensor-driven position estimation for GPS-denied environments as priority extensions. What this study demonstrates is the geometric-guidance principle itself: converting a spatial constraint from a post hoc verification metric into an active planning prior produces a framework that meets its design objectives within the tested boundaries.

Author Contributions: conceptualization and supervision, Z.T.; methodology, Z.L. and Z.T.; formal analysis, Z.L. and D.J.; software, Z.L. and Y.H.; data curation and result analysis, Z.L. and Z.T.; resources, H.L., H.Q. and T.L.; writing (draft manuscript preparation), Z.L. and Z.T.; writing (review and editing), L.L. and Z.T. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the National Natural Science Foundation of China (62303077), the Sichuan Science and Technology Program (No. 2025ZNSFSC0315), the Fund of Robot Technology Used for Special Environment Key Laboratory of Sichuan Province (No. 24kftk04), the Key Laboratory

of Detection and Application of Space Effect in Southwest Sichuan at Leshan Normal University, Education Department of Sichuan Province (No. ZDXM202401001), and the Undergraduate Innovation Training Program of Chengdu University (No. S202511079045).

Data Availability Statement: The MATLAB source code and simulation data that support the findings of this study are available from the corresponding author upon reasonable request. The Two-Line Element (TLE) orbital data used for satellite propagation are publicly accessible from Celestrak (<https://celestrak.org>).

Conflicts of Interest: Author Tao Liu was employed by Chengdu Kinyea Technologies Co., Ltd., Chengdu, China. The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

AER	Azimuth, Elevation, Range
ECEF	Earth-Centered Earth-Fixed
ECI	Earth-Centered Inertial
ENU	East–North–Up
G-U-S	Ground station–UAV–Satellite
LEO	Low Earth Orbit
LLH	Latitude, Longitude, Height
QP	Quadratic Programming
RRT*	Rapidly-exploring Random Tree Star
SGP4	Simplified General Perturbations 4
TLE	Two-Line Element
UAV	Unmanned Aerial Vehicle

References

1. Zheng, X.; Wu, Y.; Fan, L.; Lei, X.; Hu, Q.; Karagiannidis, G.K. Dual-Functional UAV-Empowered Space-Air-Ground Networks: Joint Communication and Sensing. *IEEE J. Sel. Areas Commun.* **2024**, *42*, 3412–3427. <https://doi.org/10.1109/jsac.2024.3459079>.
2. Karaman, S.; Frazzoli, E. Sampling-based Algorithms for Optimal Motion Planning. *Int. J. Robot. Res.* **2011**, *30*, 846–894. <https://doi.org/10.1177/0278364911406761>.
3. Vallado, D.A.; Crawford, P.; Hujak, R.; Kelso, T.S. Revisiting Spacetrack Report #3. In Proceedings of the AIAA/AAS Astrodynamics Specialist Conference, Keystone, CO, USA, 21–24 August 2006; AIAA 2006-6753. <https://doi.org/10.2514/6.2006-6753>.
4. Kelso, T.S. Validation of SGP4 and IS-GPS-200D Against GPS Precise Ephemerides. In Proceedings of the 17th AAS/AIAA Space Flight Mechanics Conference, Sedona, AZ, USA, 28 January–1 February 2007; AAS 07-127.
5. Montenbruck, O.; Gill, E. *Satellite Orbits: Models, Methods and Applications*; Springer: Berlin/Heidelberg, Germany, 2000. <https://doi.org/10.1007/978-3-642-58351-3>.
6. Vallado, D.A. *Fundamentals of Astrodynamics and Applications*, 5th ed.; Microcosm Press: Hawthorne, CA, USA, 2022.
7. Kassas, Z.M.; Kozhaya, S.; Kanj, H.; Saroufim, J.; Hayek, S.W. LEO Satellite Orbit Prediction via Closed-Loop Machine Learning With Application to Opportunistic Navigation. *IEEE Aerosp. Electron. Syst. Mag.* **2024**, *39*, 22–36. <https://doi.org/10.1109/MAES.2024.3432105>.
8. Kralfallah, M.; Saif, A.W.; Al-Dhafallah, M. Optimizing the Deployment of Ground Tracking Stations for Low Earth Orbit Satellite Constellations Based on Evolutionary Algorithms. *Remote Sens.* **2024**, *16*, 810. <https://doi.org/10.3390/rs16050810>.
9. Ghambari, S.; Golabi, M.; Jourdan, L.; Lepagnot, J.; Idoumghar, L. UAV Path Planning Techniques: A Survey. *RAIRO-Oper. Res.* **2024**, *58*, 2951–2989. <https://doi.org/10.1051/ro/2024073>.
10. Zeng, Y.; Zhang, R.; Lim, T.J. Wireless Communications with Unmanned Aerial Vehicles: Opportunities and Challenges. *IEEE Commun. Mag.* **2016**, *54*, 36–42. <https://doi.org/10.1109/MCOM.2016.7470933>.
11. Zeng, Y.; Xu, J.; Zhang, R. Energy Minimization for Wireless Communication with Rotary-Wing UAV. *IEEE Trans. Wirel. Commun.* **2019**, *18*, 2329–2345. <https://doi.org/10.1109/TWC.2019.2902559>.
12. Wu, Q.; Zeng, Y.; Zhang, R. Joint Trajectory and Communication Design for Multi-UAV Enabled Wireless Networks. *IEEE Trans. Wirel. Commun.* **2018**, *17*, 2109–2121. <https://doi.org/10.1109/TWC.2017.2789293>.

13. Li, B.; Fei, Z.; Zhang, Y. UAV Communications for 5G and Beyond: Recent Advances and Future Trends. *IEEE Internet Things J.* **2019**, *6*, 2241–2263. <https://doi.org/10.1109/JIOT.2018.2887086>.
14. Gammell, J.D.; Srinivasa, S.S.; Barfoot, T.D. Informed RRT*: Optimal Sampling-Based Path Planning Focused via Direct Sampling of an Admissible Ellipsoidal Heuristic. In Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Chicago, IL, USA, 14–18 September 2014; pp. 2997–3004. <https://doi.org/10.1109/IROS.2014.6942976>.
15. Wu, T.; Zhang, Z.; Jing, F.; Gao, M. A Dynamic Path Planning Method for UAVs Based on Improved Informed-RRT* Fused Dynamic Windows. *Drones* **2024**, *8*, 539. <https://doi.org/10.3390/drones8100539>.
16. Fan, J.; Chen, X.; Liang, X. UAV Trajectory Planning Based on Bi-Directional APF-RRT* Algorithm with Goal-Biased. *Expert Syst. Appl.* **2023**, *213*, 119137. <https://doi.org/10.1016/j.eswa.2022.119137>.
17. Babel, L. Online Flight Path Planning with Flight Time Constraints for Fixed-Wing UAVs in Dynamic Environments. *Int. J. Intell. Unmanned Syst.* **2022**, *10*, 416–443. <https://doi.org/10.1108/IJIUS-11-2020-0063>.
18. Machmudah, A.; Parman, S.; Abbasi, A.; Solomin, E. Flight Trajectories Optimization of Fixed-Wing UAV by Bank-Turn Mechanism. *Drones* **2022**, *6*, 69. <https://doi.org/10.3390/drones6030069>.
19. Burzyński, W.; Stecz, W. Trajectory Planning with Multiplatform Spacetime RRT*. *Appl. Intell.* **2024**, *54*, 9524–9541. <https://doi.org/10.1007/s10489-024-05650-4>.
20. Motamedi, M.; Mortazavi, A.; Sabzehparvar, M.; Duchon, F. Minimum Time Search Using Ant Colony Optimization for Multiple Fixed-Wing UAVs in Dynamic Environments. *Appl. Soft Comput.* **2024**, *165*, 112025. <https://doi.org/10.1016/j.asoc.2024.112025>.
21. Mellinger, D.; Kumar, V. Minimum Snap Trajectory Generation and Control for Quadrotors. In Proceedings of the 2011 IEEE International Conference on Robotics and Automation, Shanghai, China, 9–13 May 2011; pp. 2520–2525. <https://doi.org/10.1109/ICRA.2011.5980409>.
22. Richter, C.; Bry, A.; Roy, N. Polynomial Trajectory Planning for Aggressive Quadrotor Flight in Dense Indoor Environments. In *Robotics Research*; Inaba, M., Corke, P., Eds.; Springer Tracts in Advanced Robotics; Springer: Cham, Switzerland, 2016; Volume 114, pp. 649–666. https://doi.org/10.1007/978-3-319-28872-7_37.
23. Zhang, W.; Yadmellat, P.; Gao, Z. A Sufficient Condition for Convex Hull Property in General Convex Spatio-Temporal Corridors. In Proceedings of the IEEE Intelligent Vehicles Symposium (IV), Aachen, Germany, 4–9 June 2022; pp. 1084–1091. <https://doi.org/10.1109/IV51971.2022.9827156>.
24. Liu, S.; Watterson, M.; Mohta, K.; Sun, K.; Bhattacharya, S.; Taylor, C.J.; Kumar, V. Planning Dynamically Feasible Trajectories for Quadrotors Using Safe Flight Corridors in 3-D Complex Environments. *IEEE Robot. Autom. Lett.* **2017**, *2*, 1688–1695. <https://doi.org/10.1109/LRA.2017.2663526>.
25. Arshad, M.A.; Ahmed, J.; Bang, H. Quadrotor Path Planning and Polynomial Trajectory Generation Using Quadratic Programming for Indoor Environments. *Drones* **2023**, *7*, 122. <https://doi.org/10.3390/drones7020122>.
26. Sun, S.; Wang, H.; Xu, Y.; Wang, T.; Liu, R.; Chen, W. A Fusion Approach for UAV Onboard Flight Trajectory Management and Decision Making Based on the Combination of Enhanced A* Algorithm and Quadratic Programming. *Drones* **2024**, *8*, 254. <https://doi.org/10.3390/drones8060254>.
27. Zhang, H.; Xi, S.; Jiang, H.; Shen, Q.; Shang, B.; Wang, J. Resource Allocation and Offloading Strategy for UAV-Assisted LEO Satellite Edge Computing. *Drones* **2023**, *7*, 383. <https://doi.org/10.3390/drones7060383>.
28. Lee, Y.; Seol, K.; Song, S.; Park, H.; Park, L. UAV-Satellite Networks: A Brief Overview. In Proceedings of the 2024 15th International Conference on Information and Communication Technology Convergence (ICTC), Jeju Island, Republic of Korea, 16–18 October 2024; pp. 736–737. <https://doi.org/10.1109/ICTC62082.2024.10826759>.
29. Mohamed, E.M.; Alnakhlhi, M.; Fouda, M.M. Joint UAV Trajectory Planning and LEO-Sat Selection in SAGIN. *IEEE Open J. Commun. Soc.* **2024**, *5*, 1624–1638. <https://doi.org/10.1109/OJCOMS.2024.3372551>.
30. Chai, F.; Zhang, Q.; Yao, H.; Xin, X.; Gao, R.; Guizani, M. Joint Multi-Task Offloading and Resource Allocation for Mobile Edge Computing Systems in Satellite IoT. *IEEE Trans. Veh. Technol.* **2023**, *72*, 7783–7795. <https://doi.org/10.1109/TVT.2023.3238771>.
31. Amponis, G.; Lagkas, T.; Bouzinis, P.; Sarigiannidis, P.; Moscholios, I. Introducing Energy Efficient Routing in UAV-Satellite NTNs for Dynamic 6G Interconnectivity. *IEEE Trans. Commun.* **2025**, *74*, 1050–1065. <https://doi.org/10.1109/TCOMM.2025.3634253>.
32. Mohapatra, H. Adaptive Ant Colony Methods for UAV-LEO Coordination in Non-Terrestrial IoT. *Front. Commun. Netw.* **2025**, *6*, 1691346. <https://doi.org/10.3389/frcmn.2025.1691346>.
33. Cui, Y.; Zhang, Q.; Feng, Z.; Wen, Q.; Zhou, Y.; Wei, Z.; Zhang, P. Sensing-Assisted Accurate and Fast Beam Management for Cellular-Connected mmWave UAV Network. *China Commun.* **2024**, *21*, 271–289. <https://doi.org/10.23919/JCC.ea.2023-0140.2024.01>.
34. Zhao, J.; Gao, F.; Wu, Q.; Jin, S.; Wu, Y.; Jia, W. Beam Tracking for UAV Mounted SatCom on-the-Move With Massive Antenna Array. *IEEE J. Sel. Areas Commun.* **2018**, *36*, 363–375. <https://doi.org/10.1109/JSAC.2018.2804239>.
35. Lukito, W.D.; Xiang, W.; Lai, P.; Cheng, P.; Liu, C.; Yu, K.; Zhu, X. Integrated STAR-RIS and UAV for Satellite IoT Communications: An Energy-Efficient Approach. *IEEE Internet Things J.* **2025**, *12*, 11356–11371. <https://doi.org/10.1109/JIOT.2024.3472019>.

36. Tesfaw, B.A.; Juang, R.-T.; Tarekegn, G.B.; Kabore, W.N.; Tsai, M.-C. Joint UAV 3-D Trajectory and Resource Allocation for Integrated LEO Satellite and Multi-UAV-Enabled Marine IoT Networks: A Federated Multiagent Deep Reinforcement Learning Approach. *IEEE Internet Things J.* **2025**, *12*, 45076–45093. <https://doi.org/10.1109/JIOT.2025.3598332>.
37. Yao, W.; Chen, X.; Wang, Q.; Peng, X. Design of RIS-UAV-Assisted LEO Satellite Constellation Communication. *IEEE Trans. Commun.* **2025**, *73*, 15656–15671. <https://ieeexplore.ieee.org/document/11155890>.
38. Yu, J.; Cai, Y.; Yan, S.; Li, Y.; Wang, J.; Liu, J.; An, J. Joint 3D Beamforming-and-Trajectory Design for UAV-Satellite Uplink Covert Communication. *IEEE Trans. Commun.* **2025**, *73*, 3469–3481. <https://ieeexplore.ieee.org/document/10718357>.
39. Liang, T.; Ding, H.; Ping, Y.; Zhang, T.; Zhou, L.; Zhang, Q.; Quek, T.Q.S. Satellite-Assisted UAV Control: Sensing and Communication Scheduling for Energy-Efficient Data Collection. *IEEE Internet Things J.* **2026**, *13*, 5694–5707. <https://doi.org/10.1109/JIOT.2025.3613735>.
40. He, J.; Shi, S.; Xue, J.; Wei, S. Resource Allocation and Trajectory Optimization for NOMA-Based and UAV-Assisted Satellite IoT. *Wirel. Netw.* **2026**, *32*, 59–70. <https://doi.org/10.1007/s11276-025-04044-9>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.