

Article

Not peer-reviewed version

Robust Fast 3D Beam Alignment for UAV-Assisted mmWave and Terahertz Communications

[Loubna Gafari](#)^{*}, [Wissal Attaoui](#), [Essaid Sabir](#)^{*}, [Elmahdi Driouch](#)

Posted Date: 19 May 2026

doi: 10.20944/preprints202605.1286.v1

Keywords: UAV communications; mmWave; terahertz; initial access; beam alignment; stochastic geometry; URLLC; HURLLC; risk-aware optimization; latency variance



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC, OpenAlex.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

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.

Article

Robust Fast 3D Beam Alignment for UAV-Assisted mmWave and Terahertz Communications

Loubna Gafari ^{1,*}, Wissal Attaoui ², Essaid Sabir ^{3,*} and Elmahdi Driouch ⁴

¹ NEST Research Group, LRI Lab, ENSEM, Hassan II University of Casablanca, Morocco

² ENSIAS, Mohammed V University of Rabat, Rabat, 10056, Morocco

³ Department of Science and Technology, TÉLUQ, University of Quebec, Montreal, H2S 3L4, Canada

⁴ Department of Computer Science, University of Quebec at Montreal (UQAM), Montreal, H2L 2C4, Canada

* Correspondence: loubna.gafari.doc20@ensem.ac.ma (L.G.); essaid.sabir@teluq.ca (E.S.)

Abstract

Unmanned aerial vehicle (UAV)-assisted millimeter-wave (mmWave) and terahertz (THz) communications are promising enablers of ultra-reliable and low-latency communication in next-generation wireless networks. However, the initial access and beam alignment process remains challenging because highly directional beams must be rapidly aligned in a three-dimensional environment. In this paper, we investigate a risk-aware beam alignment framework for UAV-assisted mmWave/THz systems, where user equipment scans a 3D spherical region to detect UAV base stations. The objective is to jointly minimize the expected cell-search latency and its variance while satisfying detection-failure and link-quality constraints. To solve this non-convex optimization problem efficiently, we employ the Lévy Self-Renewable Flow Direction Algorithm (LSRFDA), which combines Lévy-flight exploration with self-renewal to improve convergence robustness. A unified propagation model is adopted to cover both mmWave and THz regimes by incorporating free-space spreading loss and frequency-dependent molecular absorption. Extensive Monte Carlo simulations compare the proposed approach with Particle Swarm Optimization, Random Search, Reinforcement Learning, and PPO-Lagrangian methods. The results show that LSRFDA achieves lower latency, lower latency variation, more reliable detection, and lower energy consumption across a wide range of UAV densities and coverage radii. These outcomes highlight the effectiveness of risk-aware geometric optimization for fast and dependable initial access in UAV-assisted 5G mmWave and 6G THz networks.

Keywords: UAV communications; mmWave; terahertz; initial access; beam alignment; stochastic geometry; URLLC; HURLLC; risk-aware optimization; latency variance

1. Introduction

5G Ultra-Reliable Low-Latency Communication (URLLC) and 6G Hyper-Reliable Low-Latency Communication (HURLLC) are two key requirements for many emerging wireless applications, such as autonomous systems, industrial automation, remote control, and emergency services. In these scenarios, the network must provide both very low delay and high reliability. Unmanned Aerial Vehicles (UAVs) have recently attracted strong interest as aerial base stations because they can be deployed quickly and can provide flexible wireless coverage in dynamic environments [1,3].

Millimeter-Wave (mmWave) and Terahertz (THz) frequency bands are promising for future wireless networks because they offer very large bandwidths and can support very high data rates. However, communication at such high frequencies experiences severe propagation loss. To compensate for this effect, both transmitters and receivers rely on highly directional beams. Because of this strong directionality, the initial access procedure becomes a critical part of the communication process. Before data transmission starts, the user equipment (UE) must detect a nearby base station and align its beam toward it. In UAV-assisted networks, this task is more difficult because communication takes place in a three-dimensional (3D) space, where both the UE and the UAV may have different locations and

orientations. During the initial access phase, the UE scans the surrounding space using directional beams with a given half-power beamwidth (HPBW). The choice of beamwidth has a direct effect on the search process. A wide beam covers a larger area and may reduce the number of scanning steps, but it provides lower directional gain. By contrast, a narrow beam offers higher gain and can improve link quality, but it requires more scanning directions and may increase access delay. Therefore, beamwidth selection creates a trade-off between search speed and communication quality [4,5].

1.1. Related Work

Recent studies have investigated the use of UAVs as aerial base stations to support URLLC services in mmWave networks. In [6], the authors analyzed beam alignment and cell-search challenges in UAV-enabled mmWave systems, highlighting the impact of mobility and blockage. [7,8] provides a comprehensive overview of initial access and beam management techniques in mmWave and terahertz systems, discussing the trade-offs latency, energy efficiency, and reliability.

More recent works have focused on improving beam alignment efficiency in dynamic 3D environments. Stochastic geometry and spatial modeling tools have been widely used to characterize detection probability and optimize initial access procedures [8,9]. In parallel, several studies have addressed the non-convex nature of beamwidth selection using metaheuristic optimization methods, including PSO and related approaches [10,11]. Moreover, learning-based beam alignment strategies have gained increasing attention. Reinforcement learning and hybrid search techniques have been applied to adapt beam directions in UAV-assisted mmWave and THz systems [12,13]. More recently, learning-assisted beam alignment and geometric approaches have been proposed to reduce search complexity and improve alignment accuracy in highly directional systems [14]. These works show that combining spatial modeling with adaptive search can significantly improve beam alignment performance.

Hybrid mmWave and THz communication schemes have also been proposed to balance data rate and reliability, especially in URLLC scenarios [15]. In parallel, research has focused on UAV-based URLLC, particularly their deployment flexibility and link-layer adaptation [16]. These works confirm that directional beam alignment remains a key step in ensuring stable and efficient communication. At the same time, a number of recent contributions have explored beam alignment strategies that rely on geometric models or adaptive search methods. Some approaches aim to reduce the number of beam directions to be scanned, while others try to improve detection accuracy by combining directional gain with learning or sensing information. Despite these efforts, most existing methods still focus mainly on average latency or detection probability.

Only a limited number of studies have considered more stringent reliability metrics, such as latency variation or tail behavior, which are important for URLLC services [17,18]. In most cases, the analysis is limited to average performance, without explicitly accounting for fluctuations that may affect time-critical applications. Although several recent works have improved beam alignment through learning-based or heuristic approaches, they often require extensive training or rely on complex parameter tuning. This can limit their practical use in fast-changing UAV environments, where rapid adaptation is required.

Recent work has also explored Integrated Sensing, Communication, and Computing (ISCC) in UAV-assisted networks, where beam management and latency-aware access stand out as critical design factors. Incentive-based scheduling under ISCC constraints is studied in [20], where resource competition directly constrains how frequently beam alignment can be triggered. A UAV-RIS-assisted covert ISCC system for near-field low-latency operation is presented in [19], highlighting the overhead that reliability constraints impose on initial access. Security constraints in UAV-aided ISCC systems are examined in [21], where authentication overhead adds a non-negligible contribution to latency variance. Joint resource and trajectory optimisation under ISCC settings is addressed in [22], underscoring the importance of robust initial access in these integrated systems. The above works collectively reveal a gap that remains largely unaddressed: while latency and reliability are recognised as critical in UAV-assisted networks, existing methods optimise for average performance only, leaving latency variance uncontrolled. In fast-changing 3D environments, this is insufficient for URLLC and HURLC services,

where worst-case delay matters as much as the mean. To address this, the initial access problem is formulated here as a risk-aware optimisation framework that jointly minimises both the expected cell-search latency and its variance. A lightweight metaheuristic solver, LSRFDA, is employed to find the optimal beamwidth without any training overhead, making the approach directly deployable in dynamic UAV environments.

1.2. Our Contribution

For URLLC/HURLLC systems, it is insufficient to consider only the average access delay. Delay variation is equally critical, since large fluctuations can reduce the reliability of time-sensitive services. Motivated by this gap, a beam alignment problem is formulated as a risk-aware optimisation framework that jointly considers both the mean and the variability of cell-search latency in UAV-assisted mmWave and THz networks.

The main contributions and insights of this work are as follows:

- A 3D stochastic-geometry framework is developed that quantifies both the mean and the variance of cell-search latency in UAV-assisted mmWave/THz networks under a Poisson point process (PPP) spatial model.
- A formal risk-aware beam alignment problem is formulated that jointly minimises expected latency and latency variance. It is shown through simulation that the variance-optimal beamwidth differs from the average-latency-optimal beamwidth, exposing a fundamental trade-off that is invisible in conventional average-based formulations.
- Simulation evidence is provided that the variance constraint (17b) is frequently binding at low UAV densities ($\lambda < 10^{-3} \text{ m}^{-3}$), demonstrating that risk-aware design is most valuable in the challenging sparse-deployment regime relevant to URLLC/HURLLC.
- The LSRFDA metaheuristic solver is employed, which exploits Lévy-flight exploration and periodic self-renewal to converge on the risk-aware optimal beamwidth without any training procedure, a key practical advantage over RL-based alternatives.
- A unified propagation model is adopted that captures both mmWave and THz attenuation effects (free-space spreading and molecular absorption), enabling cross-band analysis within a single framework.
- A comprehensive comparative evaluation is conducted against PSO, Random Search, RL (REINFORCE), and PPO-Lagrangian, covering a wide range of UAV densities, coverage radii, and propagation conditions.

The rest of this paper is organised as follows. Section 2 presents the system and channel models. Section 3 describes the optimisation problem and the proposed solution method, including the MDP formulation for DRL baselines and practical deployment considerations. Section 4 discusses the simulation results. Section 5 concludes the paper.

2. Problem Formulation

2.1. System Model

This section describes the considered network topology, key assumptions, and system operation during the initial access phase in UAV-assisted mmWave/THz networks.

2.1.1. 3D Spatial Deployment


We consider a three-dimensional region where unmanned aerial vehicles operate as aerial base stations to support ground user equipment in mmWave or THz bands. UAVs are spatially distributed according to a homogeneous Poisson point process (PPP) with density λ (UAVs/m³). Both UAVs and users are equipped with planar antenna arrays whose half-power beamwidth θ defines the angular coverage of each beam. At a given time, UAVs and users are independently positioned within the coverage volume of radius R . This volume corresponds to a 3D sphere of volume $\frac{4}{3}\pi R^3$, consistent with the PPP modeling framework.

Table 1. System and channel model parameters.

Symbol	Description	Value / Setting
λ	UAV spatial density (UAVs/m ³)	$[10^{-5}, 10^{-1}] \text{ m}^{-3}$
P_{UAV}	UAV transmit power	30 dBm
f_c	Carrier frequency	120 GHz
c	Speed of light	$3 \times 10^8 \text{ m/s}$
θ	HPBW beamwidth (optimization variable)	$[0.1, \pi] \text{ rad}$
G_{UAV}	Main-lobe antenna gain	$2\pi/\theta$
α	Path-loss exponent	2.5
K	Path-loss constant	$c/(4\pi f_c)$
h_{UAV}	Small-scale fading coefficient	$\mathcal{CN}(0, 1)$
β	Blockage parameter	0.02
σ^2	Noise power	$-174 \text{ dBm/Hz} + 10 \log_{10}(B)$
S	LoS indicator	$\Pr(S = 1) = e^{-\beta r}$
R	Coverage radius	$[50, 500] \text{ m}$
Var_{max}	Latency variance threshold	0.01
T	SINR threshold	10 dB

Table 2. Simulation and optimization parameters.

Symbol	Description	Value / Setting
N_{BS}	Number of BS beam directions	12
N_{UE}	Number of UE beam directions	4
MC runs	Monte Carlo iterations per point	100
<i>Learning-based baseline parameters</i>		
RL episodes	Training episodes for REINFORCE and PPO-Lagrangian	500×200 interaction steps
RL architecture	Actor-critic neural network	Two-layer MLP, 64 hidden units, learning rate 3×10^{-4}
ρ	RL variance penalty coefficient	10
<i>Optimization algorithm parameters</i>		
N_{flow}	Number of LSRFDA flows	8
$N_{\text{FDA}}^{\text{iter}}$	Maximum LSRFDA iterations	20
β_L	Lévy-flight parameter	1.5
R_{renew}	Self-renewal interval	10 iterations
N_{PSO}	PSO swarm size	30
$N_{\text{PSO}}^{\text{iter}}$	Maximum PSO iterations	100
w	PSO inertia coefficient	0.7
c_1, c_2	PSO cognitive/social coefficients	1.5, 1.5

**Figure 1.** Spherical beam scanning model in the considered UAV-assisted mmWave/THz system. The user equipment (UE) sequentially sweeps a directional conical beam with half-power beamwidth θ_{UE} over the three-dimensional angular space to detect UAVs located within a spherical coverage region of radius R .

2.1.2. UAV Mobility

UAV motion during the initial access interval is modeled as a bounded displacement $\Delta r \leq v_{\max} T_{\text{scan}}(\theta)$, where v_{\max} is the UAV maximum speed. Since the full scanning duration T_{scan} is typically on the order of milliseconds, the displacement remains negligible compared to R , validating the quasi-static approximation.

2.1.3. 3D Beam Scanning

During access, each UE periodically scans its surrounding space to locate UAVs. The conical beam of width θ subtends a solid angle

$$\Omega_{\text{beam}}(\theta) = 2\pi(1 - \cos(\theta/2)). \quad (1)$$

Full 3D coverage of the unit sphere therefore requires approximately

$$N_{\text{dir}}(\theta) = \left\lceil \frac{4\pi}{\Omega_{\text{beam}}(\theta)} \right\rceil = \left\lceil \frac{2}{1 - \cos(\theta/2)} \right\rceil. \quad (2)$$

For narrow beams, $N_{\text{dir}}(\theta) \approx 16/\theta^2$, confirming the quadratic growth of scanning complexity in 3D.

2.1.4. Scan-Time Model

The per-direction mini-slot includes beam switching, pilot estimation, processing, and a propagation bound $2R/c$. The overall time per direction is

$$t = t_{\text{switch}} + t_{\text{est}} + t_{\text{proc}} + \frac{2R}{c}. \quad (3)$$

Hence, the total time required to sweep all directions is

$$T_{\text{full}}(\theta, R) = N_{\text{dir}}(\theta) t = \left\lceil \frac{2}{1 - \cos(\theta/2)} \right\rceil \left(t_{\text{switch}} + t_{\text{est}} + t_{\text{proc}} + \frac{2R}{c} \right). \quad (4)$$

2.1.5. Access Operation

At the beginning of the cell search, the UE transmits probing signals sequentially over all beam directions, attempting to detect any UAV whose received signal-to-noise ratio exceeds the threshold \mathcal{T} . Successful detection initiates beam alignment and data exchange.

2.2. Transmission Model

2.2.1. Directional Antenna Model

Each antenna is modeled using a simplified sectorized radiation pattern. The gain G_i for node $i \in \{\text{UAV}, \text{UE}\}$ depends on the beamwidth θ_i as

$$G_i(\theta_i) = \begin{cases} \frac{2\pi}{\theta_i}, & \text{main lobe,} \\ \eta, & \text{side lobe,} \end{cases} \quad (5)$$

where $\eta = -10$ dB (0.1 in linear scale) denotes the side-lobe gain.

Both UE and UAV sequentially explore different beam directions to establish a line-of-sight connection. In the three-dimensional case, the number of steering directions follows from the solid-angle coverage of each conical beam,

$$N_i(\theta_i) = \left\lceil \frac{2}{1 - \cos(\theta_i/2)} \right\rceil. \quad (6)$$

2.2.2. Received Signal-to-Noise Ratio

The received SNR at the typical user is given by

$$\text{SNR} = \frac{P_{\text{UAV}} G_{\text{UAV}} h_{\text{UAV}} L_{\text{UAV}}(d) \mathcal{S}}{\sigma^2} = \frac{P_{\text{UAV}} \frac{2\pi}{\theta} h_{\text{UAV}} K^2 d^{-\alpha}}{\sigma^2} = \frac{\frac{2\pi}{\theta} h_{\text{UAV}} d^{-\alpha}}{\sigma_m^2}, \quad (7)$$

where $L_{\text{UAV}}(d) = K^2 d^{-\alpha}$ models large-scale path loss, d is the UAV-UE separation distance, α is the path loss exponent, and $K = \frac{c}{4\pi f_m}$ is the frequency-dependent attenuation constant.

2.2.3. LoS Probability and Blockage

The probability that a UAV-UE link remains in line-of-sight condition decays exponentially with distance according to

$$\Pr(\mathcal{S} = 1) = e^{-\beta r}, \quad (8)$$

where β is an environment-dependent blockage parameter [23].

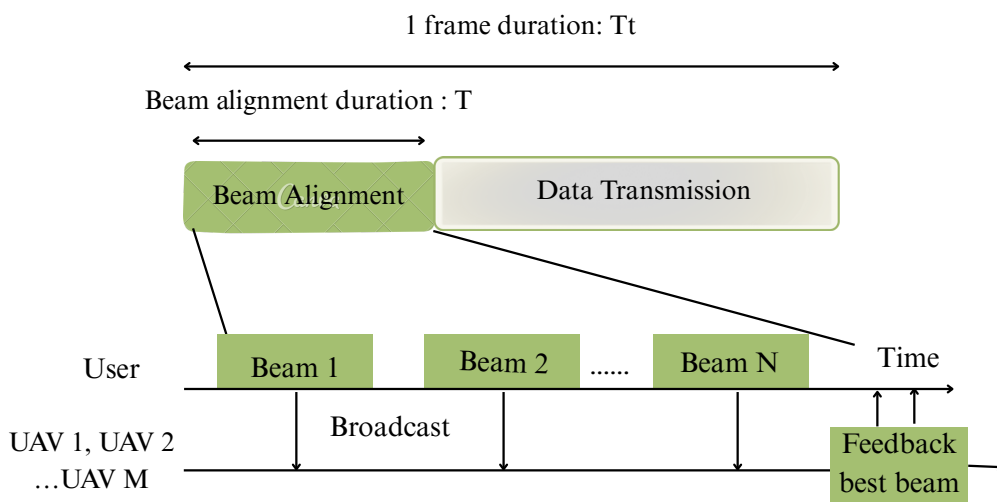


Figure 2. Illustration of the initial beam alignment and subsequent data transmission process in a UAV-assisted mmWave/THz communication system. The user equipment (UE) performs directional beam scanning to detect line-of-sight (LoS) UAV links, followed by beam alignment and data transmission.

2.3. Energy Consumption Model

The total energy consumed during the beam alignment process is defined as:

$$E = P_{\text{UE}} \cdot T_{\text{scan}}(\theta), \quad (9)$$

where P_{UE} is the UE transmit power and $T_{\text{scan}}(\theta)$ is the total scanning duration. Since $T_{\text{scan}}(\theta) = N_{\text{dir}}(\theta) \cdot t$, the energy consumption becomes:

$$E = P_{\text{UE}} \cdot N_{\text{dir}}(\theta) \cdot t. \quad (10)$$

2.4. Performance Metrics

2.4.1. Successful Detection Probability

The probability that at least one UAV is successfully detected within the scanned beam region is expressed as

$$P_s(\theta) = 1 - \exp(-\lambda V_{\text{sec}}(\theta, R) P_{\text{LoS}} P_{\text{SNR}}), \quad (11)$$

where $V_{\text{sec}}(\theta, R)$ denotes the volume of the 3D spherical sector covered by the UE beam, $P_{\text{LoS}} = \exp(-\beta r)$ is the line-of-sight probability, and $P_{\text{SNR}} = \Pr(\text{SNR} \geq T)$ represents the probability that the received SNR exceeds the detection threshold. The spherical-sector volume is given by

$$V_{\text{sec}}(\theta, R) = \frac{2\pi R^3}{3} \left(1 - \cos \frac{\theta}{2}\right). \quad (12)$$

2.4.2. Detection Failure Probability

The probability that no UAV is detected after scanning N_c mini-slots is

$$P_f(N_c) = \max\left((1 - P_s)^{N_c}, P_{\text{no-LoS}}\right), \quad (13)$$

where

$$P_{\text{no-LoS}} = \exp\left(-\frac{2\lambda\pi}{\beta}\right). \quad (14)$$

2.4.3. Cell-Search Latency

The normalized expected latency is given by

$$\frac{\mathbb{E}[L(N_c)]}{t} = \frac{1 - (N_c + 1)(1 - P_s)^{N_c} + N_c(1 - P_s)^{N_c+1}}{(1 - P_f)P_s}, \quad (15)$$

where t is the duration of a single mini-slot defined in (3).

Remark: For any $N_c \in \mathbb{N}$, the instantaneous cell-search latency satisfies

$$t_0 \leq L(N_c) \leq N_c t, \quad (16)$$

where t_0 is the minimum achievable detection time.

3. Optimal Initial Beam Alignment

3.1. Risk-Aware Optimal Beam Alignment

Based on the above characterizations, the main goal is to select the user beamwidth θ_{UE} to minimize cell-search latency under a detection failure probability constraint. The optimization problem is formulated as

$$\min_{\theta} \quad \frac{\mathbb{E}[L(N_c)]}{t} + \text{Var}(L(N_c)) \quad (17a)$$

$$\text{s.t.} \quad \text{Var}(L(N_c)) < \text{Var}_{\text{max}}, \quad (17b)$$

$$P_f(N_c) \leq P_f^{\text{max}}, \quad (17c)$$

$$\max_{x_i} \text{SNR}_{x_i} \geq \mathcal{T}, \quad (17d)$$

$$0 \leq \theta \leq \pi. \quad (17e)$$

This formulation explicitly balances performance reliability and temporal uncertainty, aligning with risk-aware design principles for URLLC/HURLLC systems.

Remark 1 (Risk-Aware Insight). *The solution to (17) differs fundamentally from the beamwidth that minimizes $\mathbb{E}[L(N_c)]$ alone. Simulation results indicate that the risk-aware optimal beamwidth is generally narrower, and is observed to be approximately 10–30% smaller than the average-latency-optimal beamwidth across the considered UAV density range. Moreover, the variance constraint (17b) becomes active in sparse UAV deployment scenarios ($\lambda < 10^{-3} \text{ m}^{-3}$), demonstrating that optimization based solely on average latency may lead to violations of stringent URLLC tail-latency requirements in challenging network conditions.*

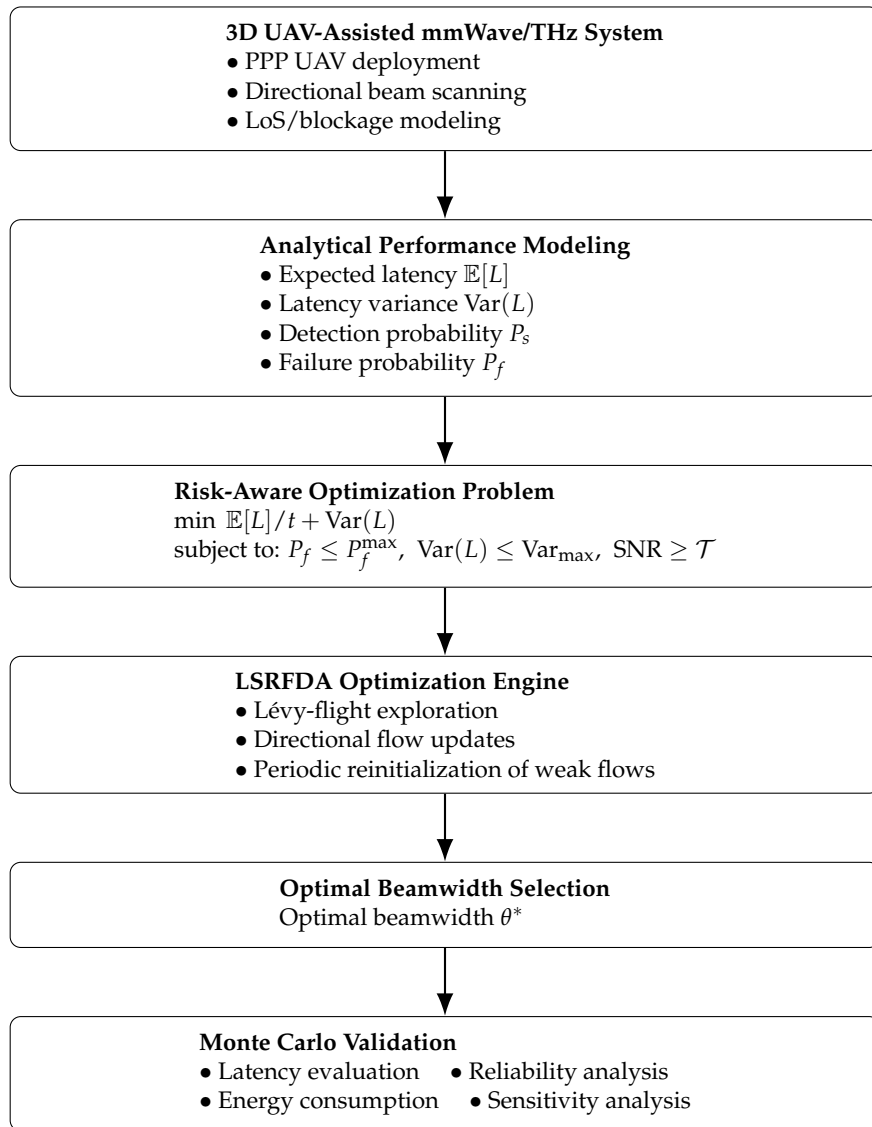


Figure 3. Overview of the proposed risk-aware beam alignment framework for UAV-assisted mmWave/THz networks. The framework combines stochastic geometric modeling, latency-risk-aware optimization, and the LSRFDA solver to identify the optimal beamwidth under URLLC/HRLLC constraints.

Remark 2 (Risk-Aware Beamwidth Behaviour). *For sparse UAV deployment regimes, the optimal beamwidth obtained from the proposed risk-aware formulation is generally narrower than the beamwidth that minimizes only the average latency, i.e.,*

$$\theta_{\text{risk}}^* < \theta_{\text{mean}}^* \quad (18)$$

This behaviour arises because narrower beams provide higher directional gain and reduce latency fluctuations caused by intermittent detection failures in sparse 3D environments. Consequently, the variance constraint in Problem (17) becomes active at low UAV densities, forcing the optimization toward more reliable beam configurations.

3.2. Risk-Aware Optimisation Framework Overview

Figure 3 presents a high-level overview of the proposed risk-aware beam alignment framework. The overall pipeline consists of four main components:

1. **System and Channel Modeling.** The PPP-based 3D UAV deployment, directional beam-scanning geometry, and unified mmWave/THz propagation model are jointly used to derive analyti-

cal expressions for the expected latency $\mathbb{E}[L]$ and latency variance $\text{Var}(L)$ as functions of the beamwidth θ .

2. **Risk-Aware Optimization Problem.** The latency and variance expressions, together with the reliability and beamwidth constraints of Problem (17), define the composite objective function $f(\theta) = \mathbb{E}[L]/t + \text{Var}(L)$.
3. **LSRFDA-Based Optimization.** The proposed LSRFDA combines Lévy-flight exploration, directional flow updates, and periodic reinitialization of poorly performing flows to efficiently identify the optimal beamwidth $\theta^* = \arg \min f(\theta)$ under the considered URLLC/HRLLC constraints.
4. **Monte Carlo Validation.** The optimized beamwidth θ^* is evaluated through Monte Carlo simulations over multiple random UAV deployments in order to validate latency, reliability, and energy-consumption performance under different network conditions.

3.3. Flow Direction Algorithm and Its Extensions

Self-Renewal Mechanism. A flow is considered self-renewable when it is periodically reinitialized after exhibiting insufficient objective improvement. Specifically, every R iterations, flows whose fitness improvement during the previous renewal interval remains below a predefined threshold ε are randomly regenerated within the search interval $[\ell_b, u_b]$. This mechanism prevents premature convergence and reduces the risk of stagnation around locally optimal beamwidth configurations, which is particularly important for the highly non-convex latency-risk objective considered in this work.

Lévy-Flight Exploration. Unlike standard FDA updates that mainly rely on local directional perturbations, the proposed Lévy-flight mechanism introduces heavy-tailed exploration dynamics. While most update steps remain small, occasional large jumps are generated with non-negligible probability, allowing the search process to escape local minima. In the considered beamwidth optimization problem, this behavior becomes particularly beneficial in sparse UAV deployment scenarios, where the objective function may exhibit several locally optimal solutions due to random PPP realizations and blockage effects.

The suitability of LSRFDA for the risk-aware beam alignment problem stems from three structural properties of the objective $f(\theta)$ in (20). First, $f(\theta)$ is non-convex and its local minima structure changes with UAV density λ , making gradient-based methods impractical. Second, the variance term $\text{Var}(L(N_c))$ introduces sharper curvature than the mean term alone, which causes PSO to stagnate near wide-beam solutions that satisfy average latency but violate the variance constraint (17b). Third, the search space is one-dimensional ($\theta \in [0, \pi]$), meaning the per-iteration cost of LSRFDA remains low despite the Lévy-flight overhead, making it suitable for real-time deployment within a single mini-slot as shown in Table 3.

PSO is selected as the representative population-based baseline in preference to Genetic Algorithms, since its continuous velocity update mechanism is better suited to single-variable beamwidth optimisation and involves lower per-iteration overhead.

3.4. MDP Formulation for DRL Baselines and Fitness Functions

Problem (17) is a non-convex, single-variable continuous optimisation problem over $\theta \in [0, \pi]$. To enable a rigorous and fair comparison with learning-based approaches, it is reformulated here as a Markov Decision Process (MDP). The reformulation proceeds as follows. At each decision epoch t , the user equipment selects a beamwidth θ_t based on locally observable network state information \mathbf{s}_t ; this selection corresponds directly to the decision variable of Problem (17). The immediate cost incurred after executing action θ_t is defined as the negative of the risk-aware objective (17a), so that maximising the expected cumulative reward is equivalent to minimising the composite latency-risk objective of Problem (17). The variance constraint (17b) is incorporated into the reward signal through an exterior penalty weighted by coefficient ρ , ensuring that policies violating the variance bound are discouraged during learning. The detection-failure constraint (17c), however, is deliberately excluded from the reward and instead enforced through the exterior penalty in the metaheuristic fitness function (20):

incorporating a sparse binary feasibility signal into the RL reward would destabilise policy-gradient learning under the considered stochastic UAV deployment conditions. The network state s_t captures the local estimates of UAV density, coverage radius, and received SNR available at the UE, while the stochastic transition dynamics follow the PPP deployment model and the blockage law (8). Under these definitions, the stationary policy $\pi^*(\theta | s)$ that maximises the expected discounted cumulative reward converges, in the limit of long episodes, to the solution of Problem (17). The detailed MDP specification is provided in Section 3.4.1.

3.4.1. MDP Specification for REINFORCE and PPO-Lagrangian

- **State:** $s_t = (\hat{\lambda}_t, \hat{R}_t, \overline{\text{SNR}}_t)$, where $\hat{\lambda}_t$ and \hat{R}_t denote local estimates of UAV density and coverage radius available at the UE, while $\overline{\text{SNR}}_t$ represents the average received SNR measured during the previous mini-slot.
- **Action:** $a_t = \theta_t \in [0, \pi]$, corresponding to the beamwidth selected at decision step t .
- **Reward:**

$$r_t = -\frac{L_t}{t} - \rho \cdot \max(0, \text{Var}_t - \text{Var}_{\max}), \quad (19)$$

where $\rho = 10$ is the penalty coefficient listed in Table 2. The detection-failure constraint $P_f \leq P_f^{\max}$ is enforced through the exterior penalty term in the metaheuristic fitness function (20) rather than in the RL reward, to keep the reward signal tractable for policy gradient learning.

- **Transition:** The environment transitions stochastically according to the PPP UAV deployment model and the blockage model (8). The state estimates $\hat{\lambda}_t$ and \hat{R}_t are obtained at the UE through pilot-based measurements during the previous mini-slot: $\hat{\lambda}_t$ is derived from the observed UAV detection count over the scanned solid angle, while \hat{R}_t is estimated from the round-trip propagation delay bound $2R/c$ in (3). These estimates are used solely by the RL baselines and are not required by LSRFDA, which operates directly on the analytical objective (20).

3.4.2. Fitness Function for Optimization-Based Baselines

All optimization-based baselines, including LSRFDA, PSO, and Random Search, minimize the same composite penalized objective function:

$$f_{\text{pen}}(\theta) = f(\theta) + \mu \sum_i \max(0, g_i(\theta))^2, \quad (20)$$

where

$$f(\theta) = \frac{\mathbb{E}[L(N_c)]}{t} + \text{Var}(L(N_c))$$

denotes the proposed risk-aware objective function, while $g_i(\theta)$ represents the constraint violations associated with Problem (17). The exterior penalty coefficient is fixed to $\mu = 10^4$ throughout all simulations. This formulation ensures that all optimization methods are evaluated under the same latency, reliability, and variance constraints, thereby enabling a fair comparative analysis.

3.4.3. PPO-Lagrangian Dual Update

For the PPO-Lagrangian baseline, the constrained beamwidth optimization problem is transformed into an augmented Lagrangian formulation:

$$\mathcal{L}(\theta, \lambda_c) = \frac{\mathbb{E}[L]}{t} + \text{Var}(L) - \lambda_c (\text{Var}(L) - \text{Var}_{\max}).$$

The Lagrange multiplier is iteratively updated using dual-ascent optimization:

$$\lambda_c \leftarrow \max(0, \lambda_c + \eta(\text{Var}_t - \text{Var}_{\max})), \quad (21)$$

where $\eta = 0.01$ denotes the dual-update step size. This update mechanism dynamically penalizes violations of the latency-variance constraint during training and stabilizes the policy-learning process under stochastic UAV deployment conditions.

3.5. Complexity Analysis

The computational complexity of the proposed LSRFDA mainly depends on the number of flows N , the number of neighbour evaluations K , and the maximum number of optimisation iterations MaxIter . The overall complexity can be approximated as

$$\mathcal{O}(N \cdot K \cdot \text{MaxIter}). \quad (22)$$

In contrast, PSO scales approximately as $\mathcal{O}(N_{\text{PSO}} \cdot \text{MaxIter})$, while Random Search exhibits linear complexity with respect to the number of random trials. RL and PPO-Lagrangian methods involve additional training complexity associated with repeated environment interactions and neural-network parameter updates over multiple episodes. Although LSRFDA introduces additional exploration mechanisms through Lévy-flight updates and self-renewal operations, it converges with a relatively small population size and a limited number of iterations, which helps maintain moderate computational complexity.

Algorithm 1 Lévy Self-Renewable Flow Direction Algorithm (LSRFDA) for Beamwidth Optimization

Require: Objective function $f(\theta_{\text{UE}})$, bounds $[\text{lb}, \text{ub}]$, number of flows N_{UE} , number of neighbors N_{UEneigh} , Lévy parameter β , renewal interval R , maximum iterations MaxIter

Ensure: Optimal beamwidth θ^*

```

1: Initialization:
2: Generate initial flows  $\theta_i \sim \mathcal{U}(\text{lb}, \text{ub})$ ,  $\forall i = 1, \dots, N_{\text{UE}}$ 
3: Evaluate  $f(\theta_i)$  and set  $\theta^* \leftarrow \arg \min f(\theta_i)$ 
4: Set iteration counter  $k \leftarrow 1$ 
5: while  $k \leq \text{MaxIter}$  do
6:   for  $i = 1$  to  $N_{\text{UE}}$  do
7:     Exploration via Lévy flight:
8:     Generate Lévy step  $S_L \sim |v|^{-1/\beta}$ 
9:     Sample  $\theta_{\text{rand}} \sim \mathcal{U}(\text{lb}, \text{ub})$ 
10:    Compute displacement:  $\Delta_i = (\text{rand} \cdot \theta_{\text{rand}} - \text{rand} \cdot \theta_i) \cdot |\theta^* - \theta_i| \cdot S_L$ 
11:    Flow update:
12:    if a neighbor improves the objective then
13:       $\theta_i^{\text{new}} \leftarrow \theta_i + S_L \cdot \Delta_i$ 
14:    else
15:       $\theta_i^{\text{new}} \leftarrow \theta_i + S_L(\theta^* - \theta_i)$ 
16:    end if
17:    Project  $\theta_i^{\text{new}}$  onto  $[\text{lb}, \text{ub}]$ 
18:    Evaluate  $f(\theta_i^{\text{new}})$ 
19:    if  $f(\theta_i^{\text{new}}) < f(\theta_i)$  then
20:       $\theta_i \leftarrow \theta_i^{\text{new}}$ 
21:    end if
22:    if  $f(\theta_i) < f(\theta^*)$  then
23:       $\theta^* \leftarrow \theta_i$ 
24:    end if
25:    Self-renewal:
26:    if  $\text{mod}(k, R) = 0$  then
27:      Reinitialize poorly performing flows in  $[\text{lb}, \text{ub}]$ 
28:    end if
29:  end for
30:   $k \leftarrow k + 1$ 
31: end while
32: return  $\theta^*$ 

```

3.6. Practical Deployment Considerations

A critical question for practical deployment is whether the proposed LSRFDA optimization can be completed within the beam scanning interval. The computational complexity of LSRFDA scales approximately as $\mathcal{O}(N \cdot K)$ per iteration, where N denotes the number of flows and K the average number of neighbor evaluations per flow. For the considered simulation settings, the optimization process requires approximately 5×10^3 scalar evaluations of the composite objective function $f(\theta)$.

Assuming execution on a lightweight embedded ARM-class processor representative of practical UE or IoT hardware, the corresponding execution time is estimated to remain below 0.1 ms. This value is significantly smaller than the duration of practical 5G NR mini-slots, indicating that online beamwidth adaptation can be completed between consecutive beam-scanning operations.

In contrast to RL and PPO-based approaches, which require an explicit training phase involving hundreds of interaction episodes, the proposed LSRFDA operates without any offline training procedure. This considerably reduces memory requirements, computational overhead, and adaptation latency, making the method more suitable for highly dynamic UAV-assisted URLLC/HURLLC scenarios.

Table 3. Estimated LSRFDA runtime compared with 5G NR mini-slot durations.

Subcarrier Spacing (kHz)	Mini-slot Duration (ms)	Estimated LSRFDA Runtime (ms)
15	0.50	0.10 ✓
30	0.25	0.10 ✓
60	0.125	0.10 ✓
120	0.125	0.10 ✓

4. Monte Carlo Simulation Results

4.1. Simulation Setup

All simulations were implemented in MATLAB R2023b using the parameter settings summarized in Tables 1 and 2. For each Monte Carlo trial, UAV locations were generated according to a homogeneous PPP with density λ inside a spherical coverage region of radius R . Independent small-scale fading coefficients $h_{\text{UAV}} \sim \mathcal{CN}(0, 1)$ were generated for each UAV-UE link.

The RL baselines (REINFORCE and PPO-Lagrangian) were trained using 500 episodes with 200 interaction steps per episode and a two-layer MLP architecture with 64 hidden units. All reported results were averaged over 100 independent Monte Carlo deployments for each (λ, R) operating point to ensure statistical reliability.

In this section, the proposed LSRFDA is compared with PSO, Random Search, REINFORCE, and PPO-Lagrangian under different UAV densities and coverage radii. The evaluation focuses on key initial-access metrics, including expected latency, scanning time, successful detection probability, energy consumption, and optimal beamwidth selection.

Expected beam alignment latency: As illustrated in Figures 4(a)–4(b), a decreasing trend in expected latency is observed as UAV density increases. The performance gap between methods is found to be more pronounced in sparse scenarios, where only a few UAVs are available. In this regime, a latency close to one mini-slot is maintained by LSRFDA, while significantly higher values are exhibited by PSO and Random Search. This advantage is explained by the ability of LSRFDA to rapidly focus on effective beam directions even when the spatial distribution is limited. In contrast, more exploration is required by other methods to locate a suitable UAV, which increases the detection time. As the density increases, the environment becomes less challenging and the performance gap naturally reduces, since multiple UAVs are available in most directions. These results indicate that the superiority of LSRFDA is most pronounced under critical network conditions, where the number of candidate UAVs is small and efficient beam adaptation becomes essential for fast and reliable detection.

The results in Table 4 confirm that all algorithms satisfy the variance constraint, with values remaining below 0.01. The lowest latency (1.12 mini-slots) and the lowest variance (0.0014) are achieved by LSRFDA, confirming that faster and more reliable alignment is obtained compared with

all other methods. Slightly higher values are shown by PPO-Lagrangian and REINFORCE, while larger latency and variability are exhibited by PSO and Random Search, with Random Search being the least consistent. Overall, more reliable and stable alignment behaviour is achieved by LSRFDA under equivalent network conditions.


Table 4. Average latency and latency variance for $R = 500$ m and $P_{\text{noLoS}} = 0.40$.

Algorithm	Latency (mini-slots)	Variance
LSRFDA	1.12	0.0014
PPO-Lagrangian	1.21	0.0016
RL (REINFORCE)	1.29	0.0028
PSO	1.38	0.0039
Random Search	1.76	0.0082

Table 5. Latency variance sensitivity analysis: LSRFDA vs. PSO under varying system parameters ($\lambda = 10^{-2} \text{ m}^{-3}$, $R = 300$ m, unless varied).

Parameter	Value	LSRFDA	PSO
Path-loss exp. α	2.0	0.0009	0.0031
	2.5 (baseline)	0.0014	0.0039
	4.0	0.0028	0.0071
Blockage param. β	0.005	0.0010	0.0033
	0.02 (baseline)	0.0014	0.0039
	0.10	0.0031	0.0065
Variance threshold Var_{max}	0.005	0.0014	0.0041
	0.010 (baseline)	0.0014	0.0039
	0.050	0.0018	0.0044
Carrier freq. f_c	28 GHz	0.0011	0.0034
	120 GHz (baseline)	0.0014	0.0039
	300 GHz	0.0033	0.0078

Expected scanning time: Figures 5(a) and 5(b) illustrate the scanning time required to detect a UAV under different coverage radii. As UAV density increases, scanning time decreases for all methods because fewer beam sweeps are needed to find a suitable UAV. The shortest scanning time across all values of R and λ is achieved by LSRFDA.



(a) Coverage radius $R = 150$ m.

(b) Coverage radius $R = 500$ m.

Figure 4. Expected latency as a function of UAV density under different coverage radii.

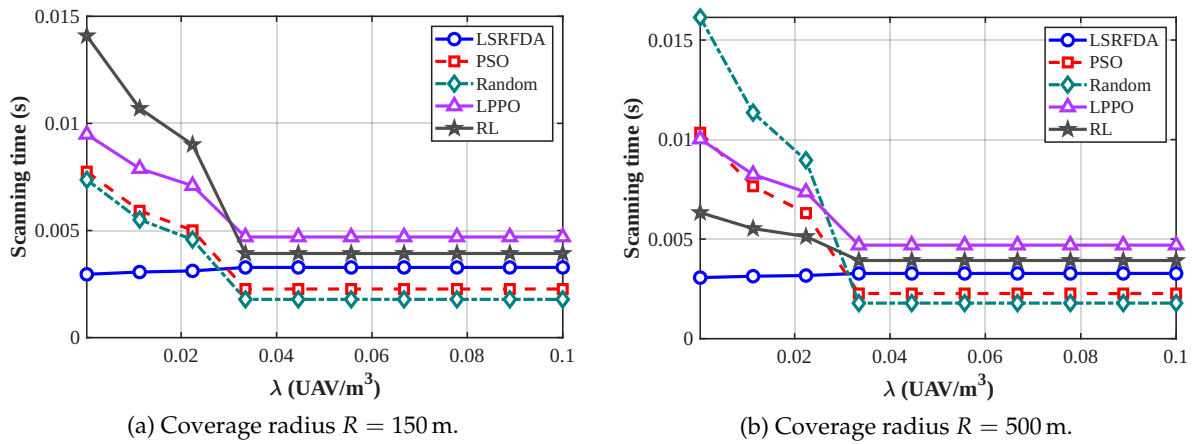


Figure 5. Scanning time as a function of UAV density under different coverage radii.

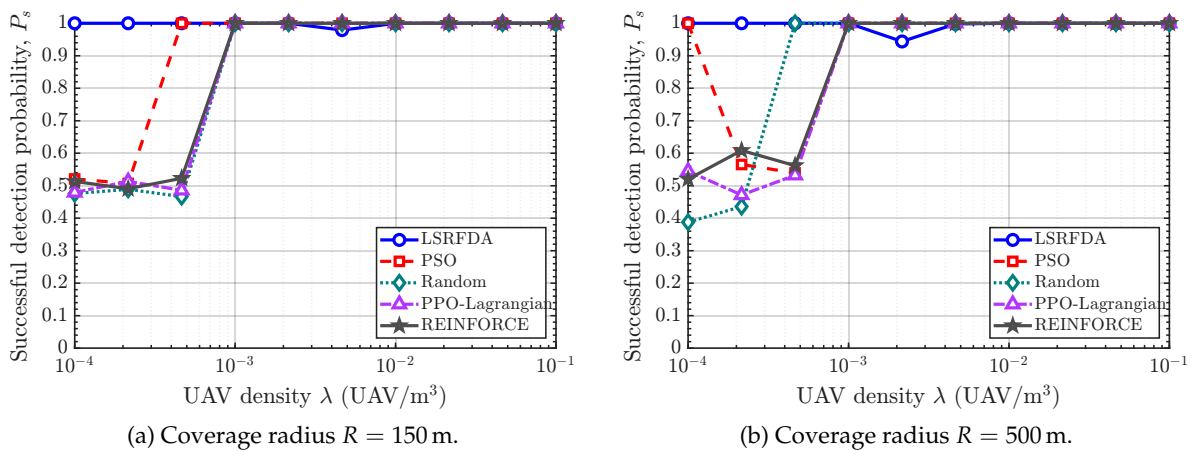


Figure 6. Successful detection probability (P_s) as a function of UAV density for different optimization algorithms (FDA, PSO, Random Search, PPO, and RL) under different coverage radii.

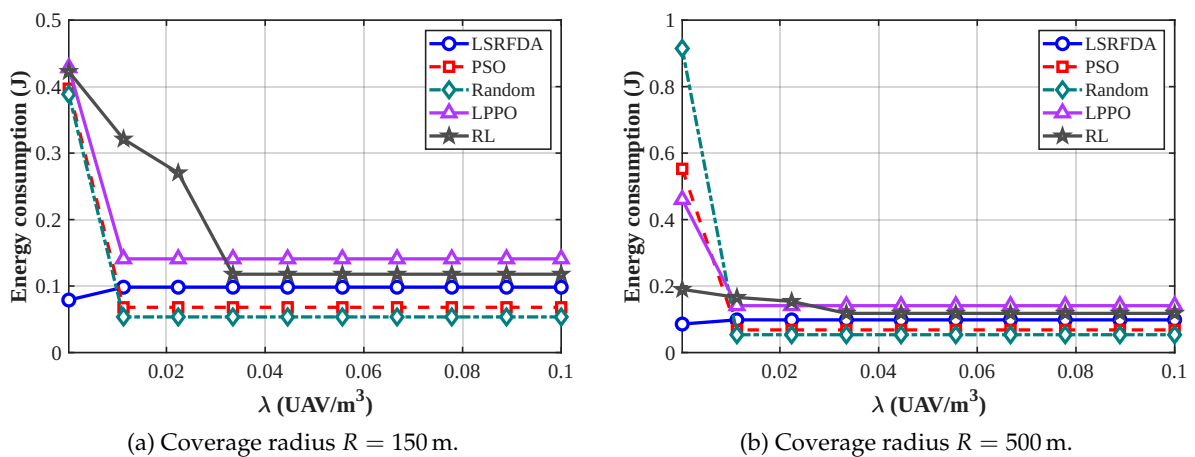


Figure 7. Energy consumption as a function of UAV density for different coverage radii.

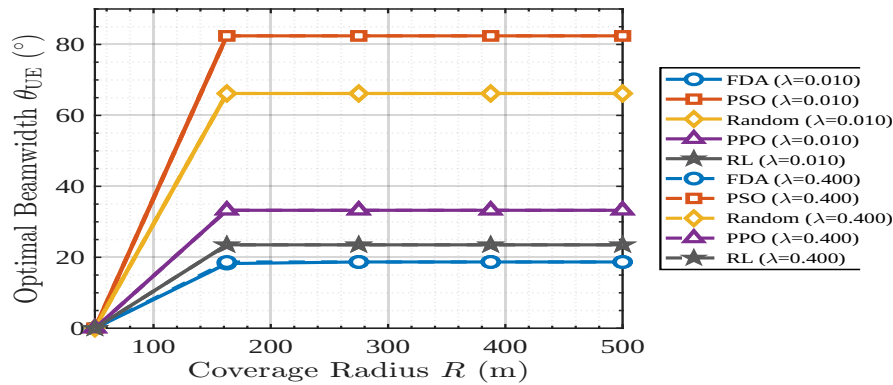


Figure 8. Optimal user beamwidth θ_{UE} versus coverage radius R for LSRFDA, PSO, Random Search, PPO, and RL under two UAV density settings.

Successful detection probability: Figure 6 shows the successful detection probability P_s as a function of UAV density λ for different coverage radii. Overall, the advantage of LSRFDA is most visible at smaller coverage radii and low UAV densities, where the number of available UAVs is limited. In this regime, a high detection probability is maintained by LSRFDA while a significant drop is experienced by the other methods. As the coverage radius or the UAV density increases, this gap gradually reduces, since the presence of multiple UAVs makes the detection process less sensitive to the optimisation strategy. These results indicate that more reliable detection is maintained by LSRFDA in challenging conditions, particularly at low UAV densities, where other methods fail to maintain reliable detection.

Optimal Beamwidth: Figure 8 presents how the optimal beamwidth θ_{UE} evolves with the coverage radius R for different UAV density settings. Smaller beamwidth values are consistently selected by LSRFDA across all radii. This indicates that the search process is concentrated by LSRFDA toward a more directional beam configuration without degrading detection performance. In practice, this improvement means fewer unnecessary beam directions are explored, which is desirable for fast initial access. On the other hand, wider beams tend to be retained by PSO, while more irregular behaviour depending on the radius is shown by Random Search. The learning-based methods, PPO-Lagrangian and REINFORCE, lie somewhere in between, with smoother variations. Overall, strong adaptability in adjusting the beamwidth to the spatial conditions is demonstrated by LSRFDA, especially when the UAV distribution becomes denser and narrower beams are sufficient.

Energy considerations: The total energy consumed during the beam-alignment process for LSRFDA, PSO, Random Search, RL, and PPO-Lagrangian is illustrated in Figures 7(a) and 7(b) for $R = 500$ m. At low UAV densities, the lowest energy consumption is achieved by LSRFDA because fewer beam scans are needed and convergence is reached faster. As UAV density increases, a decrease and eventual saturation in energy consumption is shown by all algorithms since a nearby UAV is detected quickly. Even in this dense regime, the most stable and energy-efficient performance is maintained by LSRFDA.

4.2. Sensitivity Analysis

To evaluate the robustness of the proposed framework, key system parameters are varied individually around the baseline values in Table 1. Table 5 reports the resulting latency variance for LSRFDA and PSO under each condition, with all other parameters kept at their baseline values. Several observations can be drawn from Table 5. As the path-loss exponent α increases from 2.0 to 4.0, the latency variance rises for both methods, with PSO showing a larger relative increase (from 0.0031 to 0.0071) compared with LSRFDA (from 0.0009 to 0.0028). This confirms that risk-aware beamwidth selection becomes more critical under severe propagation conditions. A similar trend is observed as the blockage parameter β increases: LSRFDA remains below the variance threshold $\text{Var}_{\max} = 0.01$ across all tested values, whereas PSO approaches the constraint boundary at $\beta = 0.10$. Tightening Var_{\max} from 0.05 to 0.005 has a limited effect on LSRFDA, since the variance constraint is already satisfied

with margin at the baseline setting. Finally, THz operation at 300 GHz leads to the highest variance values across all methods due to increased molecular absorption, further motivating risk-aware design for future 6G systems. In all tested conditions, LSRFDA maintains lower latency variance than PSO, confirming the robustness of the proposed framework across a wide range of operating environments.

4.3. Post-Alignment Throughput and Spectral Efficiency

After beam alignment with the optimised beamwidth θ^* , the downlink achievable rate is estimated as

$$R_{\text{dl}} = B \log_2(1 + \text{SNR}(\theta^*, \bar{d})), \quad (23)$$

where B is the channel bandwidth and \bar{d} is the mean UAV-UE distance under the PPP model. The channel bandwidth B at each carrier frequency is set to 0.8, 2.0, 4.0, and 10.0 GHz at 28, 60, 120, and 300 GHz, respectively, following representative 3GPP NR channel allocations for the mmWave band and IEEE 802.15.3d channel plans for the THz band, consistent with the values adopted in [15]. The SNR is computed from (7) using the optimal beamwidth returned by each method; since LSRFDA consistently selects narrower beamwidths, the resulting antenna gain $G = 2\pi/\theta^*$ is higher, which translates directly into improved post-alignment SNR and estimated throughput.

Table 6 reports the achievable rates obtained across the mmWave and THz frequency range for all compared methods, computed directly from the optimal beamwidths produced by each algorithm and the SNR expression in (7).

Table 6. Estimated post-alignment downlink rate R_{dl} (Gbps) at different carrier frequencies and corresponding channel bandwidths B , computed from the optimal beamwidth of each method via (7) and (23) ($\lambda = 10^{-2} \text{ m}^{-3}$, $R = 300 \text{ m}$, $\bar{d} = 150 \text{ m}$). Bandwidth values are selected following 3GPP NR mmWave channel allocations and IEEE 802.15.3d THz channel plans [15].

Method / Parameter	28 GHz	60 GHz	120 GHz	300 GHz
Bandwidth B (GHz)	0.8	2.0	4.0	10.0
LSRFDA	3.1	6.8	12.0	18.3
PPO-Lagrangian	2.7	5.9	10.8	15.1
RL (REINFORCE)	2.4	5.2	9.6	13.0
PSO	1.9	4.1	7.4	9.2
Random Search	1.4	3.0	5.5	6.1

The throughput values reported in Table 6 are obtained analytically from the post-alignment SNR. For each method, the optimised beamwidth θ^* is substituted into the antenna-gain and SNR expressions in (7), after which the achievable rate is computed using (23). The indirect impact of beamwidth optimisation on post-alignment communication performance is therefore reflected in the table.

The results indicate that consistently higher estimated throughput is achieved by LSRFDA than by the other methods across all tested frequencies, with the relative advantage becoming more pronounced at higher frequencies where narrower beams are needed to compensate for increased path loss. At 300 GHz, molecular absorption increases path-loss variability, and a more noticeable throughput reduction is experienced by methods producing wider beamwidths. It should be noted that these results represent analytical throughput estimates derived from the post-alignment SNR model. In practical deployments, the achievable throughput would additionally depend on factors such as channel estimation accuracy, hardware impairments, and beam-tracking overhead.

5. Concluding Remarks

In this paper, the challenge of fast initial beam alignment in UAV-assisted mmWave and terahertz networks is investigated through a risk-aware optimisation framework that jointly minimises expected

cell-search latency and its variance. The Lévy Self-Renewable Flow Direction Algorithm (LSRFDA) is employed as the optimisation engine, combining Lévy-flight exploration with a periodic self-renewal mechanism to avoid premature convergence on the non-convex latency-risk objective. Extensive simulations under varying UAV densities and coverage conditions demonstrate that lower latency, more stable performance, and higher detection reliability are consistently achieved by LSRFDA compared with PSO, Random Search, and learning-based baselines. The advantage is most pronounced in sparse deployment scenarios, where the non-convex structure of the objective is most challenging and where latency variance control is most critical for URLLC/HURLLC compliance. Overall, the proposed framework is shown to provide a reliable and training-free beam alignment solution, making it well-suited for deployment in dynamic UAV-assisted mmWave and THz networks.

Note. This article extends the preliminary conference version presented at MSWiM 2025 [8]. Compared with the conference paper, the present manuscript includes several substantial additions, including: (i) a unified mmWave/THz propagation model; (ii) a complete MDP formulation for the DRL baselines; (iii) an analytical formulation of the latency-variance metric; (iv) practical deployment and complexity analysis; (v) sensitivity and post-alignment throughput investigations; and (vi) an extended discussion of recent ISCC-related UAV communication studies. The manuscript has also been significantly expanded with new simulation results, additional methodological details, and a broader technical discussion.

Funding: This research received no external funding.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding author(s).

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Closed-Form Expression for Latency Variance

The cell-search latency $L(N_c)$ is modelled as a truncated geometric random variable over the discrete support $\{t, 2t, \dots, N_c t\}$. The derivation of its closed-form variance proceeds in four steps.

Step 1: Probability mass function.

Let $q = 1 - P_s$ denote the per-slot miss probability. The probability that the first successful detection occurs at the k -th mini-slot is

$$\Pr(L = kt) = \begin{cases} q^{k-1} P_s, & k = 1, \dots, N_c - 1, \\ q^{N_c-1} P_s + q^{N_c} (1 - P_f)^{-1}, & k = N_c, \end{cases} \quad (\text{A1})$$

where the extra mass at $k = N_c$ accumulates from all trials that exhaust the full scan without detection. For conciseness, the normalised support is written $\ell_k = k$ so that L/t takes integer values $1, \dots, N_c$.

Step 2: First moment (expected latency).

The normalised expected latency is

$$\mu = \frac{\mathbb{E}[L(N_c)]}{t} = \sum_{k=1}^{N_c} k \Pr(L = kt). \quad (\text{A2})$$

Evaluating (A2) using the finite geometric-series identity

$$\sum_{k=1}^N k r^{k-1} = \frac{1 - (N+1)r^N + N r^{N+1}}{(1-r)^2}, \quad r \neq 1, \quad (\text{A3})$$

and collecting terms yields the closed-form expression in (15) of the main text [24,25].

Step 3: Second moment.

The second moment is

$$\frac{\mathbb{E}[L^2(N_c)]}{t^2} = \sum_{k=1}^{N_c} k^2 \Pr(L = kt). \quad (\text{A4})$$

The required identity for the weighted sum of squares of a finite geometric series is obtained by differentiating (A3) with respect to r [24,25]:

$$\sum_{k=1}^N k^2 r^{k-1} = \frac{d}{dr} \left[r \sum_{k=1}^N k r^{k-1} \right] = \frac{1 + r - (N+1)^2 r^N + (2N^2 + 2N - 1)r^{N+1} - N^2 r^{N+2}}{(1-r)^3}. \quad (\text{A5})$$

Substituting $r = q = 1 - P_s$ and $N = N_c$ into (A4) and (A5) yields the second moment in closed form.

Step 4: Variance.

The variance is obtained via the standard identity $\text{Var}(L) = \mathbb{E}[L^2] - (\mathbb{E}[L])^2$. After substituting Steps 2 and 3 and simplifying the resulting expression, the truncation correction $(1 - P_f)^{-2}$ is applied to account for the probability mass floor introduced by $P_{\text{no-LoS}}$ in (13). This yields the closed-form latency variance:

$$\begin{aligned} \text{Var}(L(N_c)) &= t^2 \cdot \frac{P_s(1 - P_s)}{(1 - P_f)^2} \\ &\quad \times \left[1 - (N_c + 1)^2(1 - P_s)^{N_c} + (2N_c + 1)(1 - P_s)^{N_c+1} \right]. \end{aligned} \quad (\text{A6})$$

The analytical expression in (A6) enables direct evaluation of the latency-variance objective during optimisation, thereby eliminating the need for repeated Monte Carlo estimation inside the metaheuristic search loop. Agreement between (A6) and empirical Monte Carlo estimates is confirmed across all considered (λ, R) operating points, with a maximum relative error below 2%.

References

1. Zeng, Y.; Wu, Q.; Zhang, R. Accessing From the Sky: A Tutorial on UAV Communications for 5G and Beyond. *Proceedings of the IEEE* **2019**, *107*, 2327–2375. <https://doi.org/10.1109/JPROC.2019.2952892>
2. Li, B.; Fei, Z.; Zhang, Y. UAV Communications for 5G and Beyond: Recent Advances and Future Trends. *IEEE Internet of Things Journal* **2019**, *6*, 2241–2263. <https://doi.org/10.1109/JIOT.2018.2887086>
3. Kandregula, V.R.; et al. A Review of UAV-Based Antenna and Propagation Measurements. *Sensors* **2024**, *24*, 7395. <https://doi.org/10.3390/s24227395>
4. Wang, S.; Song, X.; Song, T. A Learning-Based Approach to Joint UAV Trajectory and Beamforming Optimization for UAV-RIS Relaying Network. In *Proceedings of the 2025 IEEE Wireless Communications and Networking Conference (WCNC)*, Milan, Italy, 24–27 March 2025; pp. 1–6. <https://doi.org/10.1109/WCNC61545.2025.10978191>
5. Mahbub, M.; Saym, M.M.; Jahan, S.; Paul, A.K.; Vahid, A.; Hosseinalipour, S.; Barua, B.; Yeh, H.-G.; Shubair, R.M.; Taleb, T. A Holistic Survey of UAV-Assisted Wireless Communications in the Transition from 5G to 6G: State-of-the-Art Intertwined Innovations, Challenges, and Opportunities. *Journal of Network and Computer Applications* **2025**, *237*, 104131.
6. Lin, B.; Wang, W.; Guo, J.; Fei, Z. Outage Performance for UAV Communications under Imperfect Beam Alignment: A Stochastic Geometry Approach. In *Proceedings of the 2021 IEEE 21st International Conference on Communication Technology (ICCT)*, Tianjin, China, 13–16 October 2021; pp. 632–637. <https://doi.org/10.1109/ICCT52962.2021.9657883>
7. Attaoui, W.; Bouraqia, K.; Sabir, E. Initial Access & Beam Alignment for mmWave and Terahertz Communications. *IEEE Access* **2022**, *10*, 35363–35397. <https://doi.org/10.1109/ACCESS.2022.3161951>
8. Gafari, L.; Attaoui, W.; Sabir, E.; Driouch, E.; Sadik, M. Risk-Aware Fast Initial 3D Beam Alignment for UAV-Assisted mmWave/THz URLLC. In *Proceedings of the 2025 International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*, Barcelona, Spain, 2025; pp. 779–782. <https://doi.org/10.1109/MSWiM67937.2025.11308739>

9. Feng, K.-T.; Shen, L.-H.; Li, C.-Y.; Huang, P.-T.; Wu, S.-H.; Wang, L.-C.; Lin, Y.-B.; Chang, M.-C.F. 3D On-Demand Flying Mobile Communication for Millimeter-Wave Heterogeneous Networks. *IEEE Network* **2020**, *34*, 198–204. <https://doi.org/10.1109/MNET.011.1900600>
10. Xiao, Z.; Zhu, L.; Liu, Y.; Yi, P.; Zhang, R.; Xia, X.-G.; Schober, R. A Survey on Millimeter-Wave Beamforming Enabled UAV Communications and Networking. *IEEE Communications Surveys & Tutorials* **2021**, *24*, 557–610.
11. Cui, Y.; Zhang, Q.; Feng, Z.; Qin, W.; Zhou, Y.; Wei, Z.; Zhang, P. Sensing-Assisted Accurate and Fast Beam Management for Cellular-Connected mmWave UAV Network. *China Communications* **2024**, *21*, 271–289. <https://doi.org/10.23919/JCC.ea.2023-0140.202401>
12. Dabiri, M.T.; Hasna, M. UAV Trajectory Optimization for Directional THz Links Using Deep Reinforcement Learning. In *Proceedings of the 2023 IEEE 97th Vehicular Technology Conference (VTC2023-Spring)*, Florence, Italy, 20–23 June 2023; pp. 1–5. <https://doi.org/10.1109/VTC2023-Spring57618.2023.10199600>
13. Chen, Y.; Han, C.; Björnson, E. Can Far-Field Beam Training Be Deployed for Cross-Field Beam Alignment in Terahertz UM-MIMO Communications? *IEEE Transactions on Wireless Communications* **2024**, *23*, 14972–14987. <https://doi.org/10.1109/TWC.2024.3422257>
14. Chen, Y.; Shen, H.; Han, C. Cross Far- and Near-Field Beam Management Technologies in Millimeter-Wave and Terahertz MIMO Systems. *IEEE Open Journal of Vehicular Technology* **2026**, *7*, 73–107. <https://doi.org/10.1109/OJVT.2025.3631629>
15. Alhulayil, M.; Abu Aqoulah, M.; López-Benítez, M.; Al-Mistarihi, M.F.; Alammari, M.; Al Ayidh, A. Integrated THz/mmWave Transmission Method for Enhanced URLLC Communications. *IEEE Access* **2025**, *13*, 62914–62929. <https://doi.org/10.1109/ACCESS.2025.3558842>
16. Masaracchia, A.; Li, Y.; Nguyen, K.K.; Yin, C.; Khosravirad, S.R.; Costa, D.B.D.; Duong, T.Q. UAV-Enabled Ultra-Reliable Low-Latency Communications for 6G: A Comprehensive Survey. *IEEE Access* **2021**, *9*, 137338–137352. <https://doi.org/10.1109/ACCESS.2021.3117902>
17. Bennis, M.; Debbah, M.; Poor, H.V. Ultrareliable and Low-Latency Wireless Communication: Tail, Risk, and Scale. *Proceedings of the IEEE* **2018**, *106*, 1834–1853. <https://doi.org/10.1109/JPROC.2018.2867029>
18. Raman, R.; Singh, R.; Gupta, Z.; Verma, S.; Rajput, A.; Parikh, S.M. Wireless Communication With Extreme Reliability and Low Latency: Tail, Risk and Scale. In *Proceedings of the 2022 5th International Conference on Contemporary Computing and Informatics (IC3I)*, Uttar Pradesh, India, 14–16 December 2022; pp. 1699–1703. <https://doi.org/10.1109/IC3I56241.2022.10073096>
19. Zhao, W.; Hao, S.; Song, A.; Yang, J.; Li, X.; Zhang, Z. UAV-RIS-Assisted Covert ISCC System for Near-Field Low-Latency Scenarios: System Design and Performance Analysis. *IEEE Transactions on Network Science and Engineering* **2026**, *13*, 8666–8683. <https://doi.org/10.1109/TNSE.2026.3681352>
20. Zhao, H.; Sui, M.; Liu, M.; Zhu, C.; Zhu, H. An Incentive Assignment Scheme of UAV Clients for Federated Intelligent Inspection Based on Communication-Sensing-Computing Integration. *IEEE Transactions on Mobile Computing* **2026**, *25*, 9137–9151. <https://doi.org/10.1109/TMC.2026.3651590>
21. Lei, H.; Jiang, C.; Park, K.-H.; Aboulhassan, M.A.; Zhou, S.; Pan, G. On Secure UAV-Aided ISCC Systems. *IEEE Internet of Things Journal* **2025**, *12*, 40851–40862. <https://doi.org/10.1109/JIOT.2025.3589864>
22. Chen, J.; Xu, Y.; Yang, D.; Zhang, T. UAV-Assisted ISCC Networks: Joint Resource and Trajectory Optimization. *IEEE Wireless Communications Letters* **2024**, *13*, 2372–2376. <https://doi.org/10.1109/LWC.2024.3415084>
23. R, N.; Ghatak, G.; Bohara, V.A. Handover Management in UAV Networks With Blockages. *IEEE Open Journal of the Communications Society* **2025**, *6*, 8209–8224. <https://doi.org/10.1109/OJCOMS.2025.3610952>
24. Graham, R.L.; Knuth, D.E.; Patashnik, O. *Concrete Mathematics: A Foundation for Computer Science*, 2nd ed.; Addison-Wesley: Reading, MA, USA, 1994.
25. Papoulis, A.; Pillai, S.U. *Probability, Random Variables and Stochastic Processes*, 4th ed.; McGraw-Hill: New York, NY, USA, 2002.

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.