



Quantum-Secured AI-Driven Drone Logistics for Real-Time Healthcare Delivery

Sana Hafeez¹ · Ghulam E Mustafa Abro² · Maran Marimuthu³

Received: 7 July 2025 / Accepted: 22 January 2026
© The Author(s) 2026

Abstract

This paper presents a quantum-resilient autonomy stack for medical-drone delivery that elevates communications and cryptography to first-class, stateful variables within motion planning rather than downstream constraints. A three-layer Air–Ground–Communications architecture integrates beyond-6 G (FR3/THz) links with ultra-reliable low-latency communication (URLLC) fallback and couples BB84-style quantum key distribution (QKD) with principled post-quantum cryptography (PQC) switching to ensure cryptographic continuity under mobility and adverse weather. At the core, a hybrid AI–RKF45 controller fuses a lightweight neural policy with the integrator’s local error, co-adapting control aggressiveness and solver step size for stiffness-aware manoeuvres and rapid re-optimisation under disturbances. The planner directly ingests link SNR, URLLC queueing delay, QKD quantum bit error rate (QBER), secure key rate (SKR), key-buffer levels, meteorological risk, and battery state-of-health, shaping a multi-objective cost that jointly minimises energy and *control-loop end-to-end latency* (sensing–compute–communication round-trip), while enforcing geofencing and Beyond Visual Line of Sight (BVLOS) constraints. In large-scale regional simulations (200 km × 200 km; up to 200 UAVs), the framework achieves a mean control-loop latency of 2.34 s (distinct from door-to-door delivery time), an 18% reduction in energy per sortie, and a 98.2% mission success rate, outperforming static planners, deep reinforcement learning (PPO/DQN), and model predictive control under matched compute budgets. The QKD↔PQC state machine applies conservative thresholds (QBER ≤ 8%, SKR ≥ 5 kbps with key-buffer hysteresis), yielding rare, millisecond-scale fallbacks that preserve latency guarantees. Complexity scales as $\mathcal{O}\left(\frac{n}{v}T\right)$, maintaining 2–5 s update times for fleets exceeding 200 vehicles. Reproducible artefacts and a staged pathway from hardware-in-the-loop to BVLOS trials support near-term healthcare deployment.

Keywords Medical drone delivery · Quantum-resilient communications · Beyond-6G networks · Autonomous UAV navigation · Control-loop latency optimisation

Maran Marimuthu and Ghulam E. Mustafa Abro have contributed equally to this work.

✉ Ghulam E Mustafa Abro
mustafa.abro@ieee.org

Sana Hafeez
S.Hafeez@ljmu.ac.uk

Maran Marimuthu
Maran.marimuthu@utp.edu.my

¹ Digital Innovation Research Institute, School of Computer Science and Mathematics, Liverpool John Moores University, Liverpool L3 3AF, UK

² Artificial Intelligence in Robotics Laboratory (AiR Lab), Electrical and Computer Engineering Department, Aarhus University, 8000 Aarhus, Aarhus Centrum, Denmark

1 Introduction

Background and Future Outlook

When care is measured in seconds, logistics becomes the lifeline. Looking towards the 2030–2050 horizon, equitable access to urgent therapies and diagnostics will increasingly depend on the safe, secure, and predictable movement of medical assets across both urban and rural corridors [1–3]. Unmanned aerial vehicles (UAVs) provide a strong complement to ground transport by bypassing congestion, terrain constraints, and disrupted infrastructure, enabling time-critical delivery of blood products, vaccines, labora-

³ Department of Management and Humanities, Universiti Teknologi PETRONAS, 32610, Perak Seri Iskandar, Malaysia



tory samples, and emergency medicines [4–6]. Achieving this capability at scale requires an autonomy stack that is robust, quantum-resilient, regulation-aware, and adaptive to environmental variability.

Limitations and Existing Work

Research into beyond-5G and beyond-6G (B6G) connectivity targets ultra-reliable low-latency communication (URLLC) with operation in Frequency Range 3 (FR3) and Terahertz (THz) bands anticipated by IMT-2030 roadmaps [7–9]. In parallel, quantum key distribution (QKD) has matured as a route to information-theoretic confidentiality, while post-quantum cryptography (PQC) provides practical, standards-track alternatives and fallbacks [10, 11]. Integration of these capabilities on agile aerial platforms remains challenging. High-frequency backhaul is strongly weather-dependent and imposes tight requirements on pointing, acquisition and tracking (PAT). Airborne QKD must maintain beam alignment under motion and manage secret-key rate (SKR) and quantum bit error rate (QBER) within power and payload budgets. The autonomy stack must react to cryptographic and spectral state as readily as it responds to wind, precipitation, temperature and geofencing [12–16].

The literature has advanced important, but largely fragmented, parts of this problem. Medical-drone pilots often adopt simplified airspace models or static communications and security regimes [17]. Deep reinforcement learning (DRL) and model predictive control (MPC) improve path planning and collision avoidance; however, typical formulations optimise kinematics and energy while treating communications and security as exogenous. Joint optimisation of latency, reliability and cryptographic continuity is rare, and most planners ignore the numerical solver's local error, a factor that matters under stiff dynamics, gusts and rapid re-tasking [18–21]. Terrestrial and satellite QKD studies clarify SKR and QBER envelopes and switching conditions but seldom address integration with low-altitude, mobile control-loop timing [10–22]. A gap therefore persists between point solutions in isolation and unified, regulation-aware stacks suited to future health corridors.

The proposed architecture is framed to be consistent with emerging guidance for BVLOS operations within unmanned traffic management (UTM) ecosystems, including altitude bands for routine medical missions, dynamic geofencing, strategic deconfliction and audit-ready flight logging. The communications stack aligns with IMT-2030 trajectories for B6G service classes and with 3GPP efforts on non-terrestrial integration, while feasibility bounds for FR3 and free-space optical links, pointing and tracking stability, and airborne QKD are stated explicitly so that near-term deployments are separated clearly from forward-looking capabilities.

Motivation

Closing this gap requires treating communications and cryptography not as fixed constraints but as stateful signals that inform motion. A healthcare-grade autonomy system should adapt its route if the signal-to-noise ratio (SNR) degrades, if URLLC queueing increases, or if QKD keys become scarce. Seamless fallback to PQC under explicit operational thresholds is necessary, and this must be achieved while maintaining energy efficiency, BVLOS compliance and robust control performance in adverse weather.

The problem considered here is real-time routing for a fleet of n UAVs delivering clinical payloads under communications and security constraints. For each drone d over horizon $t = 0, \dots, T$, the planner selects control inputs $u_d(t)$ that generate a trajectory $x_d(t)$ governed by vehicle dynamics integrated with RKF45. The objective is to minimise a composite cost

$$\min_{\{u_d(t)\}} \sum_{d=1}^n \sum_{t=0}^T [w_d L_d(t) + w_e E_d(t) + w_s \Psi(\text{QBER}(t), \text{SKR}(t))]$$

subject to geofencing and collision-avoidance constraints, communications feasibility $\text{SNR}(t) \geq \text{SNR}_{\text{th}}$, cryptographic continuity $\text{SKR}(t) \geq R_{\text{QKD}}^{\min}$ or PQC fallback, and regulatory altitude and BVLOS limits. The function $\Psi(\cdot)$ penalises quantum-security risk and triggers QKD to PQC switching when $\text{QBER}(t) > \text{QBER}_{\text{max}}$. The term control-loop end-to-end latency denotes the sensing to compute to link round-trip that governs flight safety and is reported in seconds. Mission completion time denotes door-to-door delivery time and is reported in minutes; the two measures are not interchangeable.

Approach in Brief

An end-to-end, security and communications-aware autonomy framework for UAV healthcare logistics is proposed. The architecture adopts B6G-class connectivity for control-plane reliability, couples QKD with a PQC fallback to ensure cryptographic continuity, and embeds these signals within the planner so that trajectories adapt to wireless conditions, key availability and regulatory constraints alongside environmental variables. A key technical pillar is a hybrid controller that couples a lightweight neural policy with RKF45 integration. Unlike DRL-only or MPC-only pipelines, the controller ingests a normalised proxy of the integrator's local truncation error and uses it to modulate both action aggressiveness and integration step size on the fly. The result is stiffness-aware manoeuvres and rapid re-optimisation under disturbance, with transparent and auditable dynamics [19–21].

Novelty and Contributions

The central innovation lies in coupling a lightweight neural policy with the integrator's own error feedback so that action aggressiveness and integration step size co-adapt to stiffness, disturbance and re-tasking, which conventional deep reinforcement learning and model predictive control pipelines do not exploit [27]. A second innovation is to treat communications and cryptography as stateful signals within the planner rather than as downstream constraints. Quantum bit error rate and QKD, together with signal-to-noise ratio and ultra-reliable low-latency queueing from the B6G link, shape the objective and trigger cryptographic continuity through principled switching between QKD and PQC. This combination yields a unified, regulation-aware stack for UAV healthcare logistics that integrates solver-aware control with security and communications co-optimisation.

Medical Drone Corridors for the Mid-Century

Medical drone corridors during the 2030s are likely to be supported by heterogeneous networks that treat airspace as a first-class service domain. B6G backbones across FR2 and FR3, complemented by free-space optical links, will coexist with resilient sub-6GHz overlays, while non-terrestrial networks such as high-altitude platforms and low-Earth-orbit constellations provide ubiquitous control-plane reach and disaster continuity. These layers will be orchestrated by network digital twins that forecast congestion, spectrum occupancy and micro-weather, then pre-position capacity, keys and compute. Joint communication and sensing will make the link itself a sensor, with beams measuring aerosol density and turbulence along the path, and reconfigurable surfaces on rooftops shaping propagation to enforce geofencing. Security will evolve towards continuity of trust, with hybrid QKD and PQC fabrics allocating keys as a schedulable resource and lightweight attestation ensuring drones remain cryptographically present in the corridor. Computation will follow the patient through edge resources at hospitals, ambulances and high-altitude platforms, with semantic communication prioritising task-level intents over raw data to deliver reliability where it carries most clinical value. Governance will mature into evidence-carrying automation so that UTM can interoperate with hospital command systems, missions can publish cryptographic flight logs suitable for audit without exposing identity, and corridors can expose live reliability, coverage and key-continuity scores with confidence bounds. The eventual measure of success will be a utility-like fabric for urgency that commits not only a path and a vehicle, but the sensing, keys, compute and energy required to deliver care with proof when seconds decide outcomes.

Paper Organisation and Metric Conventions

Section II develops the system architecture across air, ground and communications layers, including B6G and THz feasibility bounds and BVLOS compliance. Section III details the methodology: the hybrid AI-RKF45 planner, the quantum-resilient security model with QKD to PQC switching, and the integration of communications and cryptographic state within the optimisation. Section IV sets out the simulation methodology and baselines and reports results across latency, energy, reliability and security metrics with appropriate confidence intervals. Section V discusses practical feasibility, limitations and a staged roadmap from simulation to hardware-in-the-loop and field trials. Table 1 summarises notation and acronyms used throughout. For clarity, control-loop end-to-end latency, the sensing to compute to link round-trip relevant to flight safety, is distinguished from door-to-door mission completion time.

2 System Architecture and Problem Formulation

This section formalises a three-layer architecture for B6G healthcare logistics using autonomous uncrewed aerial vehicles (UAVs). The stack comprises an *Air Layer* for flight dynamics and energy constraints, a *Ground Layer* for healthcare facilities and demand points, and a *Communications Layer* for secure, high-throughput networking that integrates QKD with PQC fallback. The optimisation objective is to minimise *control-loop end-to-end latency* and energy usage while maintaining reliability and *cryptographic continuity* under adverse weather and regulatory constraints.

2.1 Operational Scenario and Environment

The target environment spans mixed urban and rural regions, including mountainous and coastal terrain as well as island communities. High-resolution digital elevation models (DEMs) at 1 m spatial resolution capture topographical gradients and obstacle envelopes. Wind fields, precipitation and temperature are modelled as slowly varying exogenous processes that drive a composite weather index $WI \in [0, 1]$ used for rate adaptation and replanning (Sect. 2.8). Multi-year healthcare demand data (aggregated records for medicines, blood products and time-critical pathology) exhibit diurnal and seasonal surges that shape spatio-temporal load. The simulated population coverage may reach 500,000 residents across metropolitan centres and remote regions.

All simulations run on a high-performance computing cluster with dual Intel Xeon processors (up to 80 logical cores), 512GB ECC RAM, and NVIDIA A100 GPUs (80GB HBM2). This platform enables sub-second control-



Table 1 Key Parameters, Notations, and Acronyms

Acronym	Description	Acronym	Description
Ω	3D UAV operational space (\mathbb{R}^3)	T	Discrete time horizon $[0, T_{\max}]$
\mathcal{L}_A	Air Layer: UAV mobility and logistics	N_D	Number of autonomous drones
$\mathbf{S}_d(t)$	State vector of drone d at time t	(x_d, y_d, z_d)	3D coordinates of drone d
$v_d(t)$	Velocity of drone d at time t	$E_d(t)$	Remaining battery energy at time t
$P_d(t)$	Payload weight carried by drone d	$\theta_d(t)$	Pitch angle (altitude control)
$\psi_d(t)$	Yaw angle (horizontal control)	v_{\min}, v_{\max}	Minimum/maximum velocity bounds
h_{\min}, h_{\max}	Minimum/maximum altitude bounds	$\dot{E}_d(t)$	Energy consumption rate
ω_d	Delay weighting factor for drone d	P_{hover}	Hovering power consumption
$P_{\text{flight}}(v_d)$	Motion power at velocity v_d	P_0, P_1, P_2	Aerodynamic power coefficients
E_{\min}	Minimum energy for safe return-to-base (RTB)	\mathcal{L}_G	Ground Layer: healthcare facilities
$L_h(i)$	State of healthcare facility i	$L_v(j)$	State of rural site j
$\phi_h(i), \lambda_h(i)$	Coordinates of healthcare facility i	$C_h(i)$	Medical supply capacity of facility i
$\mu_h(i)$	Medical demand at facility i	$\phi_v(j), \lambda_v(j)$	Coordinates of rural site j
$p_v(j)$	Population density at site j	$d_v(j)$	Medical demand at site j
\mathcal{L}_C	Communication Layer: THz and QKD links	$h_{d,g}(t)$	Channel gain between drone d and ground station g
$\beta_{d,g}(t)$	Large-scale fading factor	$PL_{d,g}$	Path loss between drone and ground
$d_{d,g}$	Distance between drone and ground node	f_c	Carrier frequency (THz band)
$\zeta(\omega, h)$	Atmospheric attenuation function	X_σ	Log-normal shadowing component
R_{QKD}	Secure key rate in QKD	$QBER_{\text{threshold}}$	Maximum acceptable QBER _{th}
SNR_{th}	Minimum signal-to-noise ratio for secure link	W_d	Latency weighting factor
W_e	Energy efficiency weighting factor	W_s	QKD security weighting factor
η_{SNSPD}	Detection efficiency of SNSPDs	R_{QKD}^{\min}	Minimum secure key rate threshold
P_{QKD}	Power consumption for QKD operations	$L_d(t)$	Latency for drone d at time t
$R_d(t)$	Reliability metric for drone d at time t	\hat{G}	Estimated trajectory gradient
RKF45	Runge–Kutta–Fehlberg 45 integration method	n	Number of drones in simulation
T_{\max}	Maximum flight duration	ν	Number of parallel compute cores
$\mathcal{O}(nT)$	Computational complexity (serial)	$\mathcal{O}(n/\nu \cdot T)$	Complexity after parallelisation
W_v	Wind velocity component	W_p	Precipitation intensity
W_t	Temperature variation	WI	Composite weather index

loop intervals and parallel evaluation of up to 200 UAVs over domains as large as $200 \text{ km} \times 200 \text{ km}$. Figures are illustrative; the methods and parameters are reported in Sect. 4.

2.2 Three-Layer Architecture

Figure 1 presents the architecture. The *Air Layer* models multirotor and fixed-wing UAVs with flight envelopes and battery ageing. The *Ground Layer* comprises hospitals, clinics and remote communities with time-varying demand and geofencing/UTM constraints. The *Communications Layer* provides B6G-class URLLC links in FR3 and free-space optical (FSO) bands and integrates BB84-style QKD for quantum-resilient keying. QKD continuity is maintained by a PQC fallback when operational thresholds are breached (Sect. 2.7). Figure 2 illustrates an example network snapshot with throughput-coded nodes and proximity edges.

2.3 Layer–Objective Mapping

To make explicit how technology choices drive system objectives, Table 2 maps each layer to the optimisation targets: latency L , energy E and reliability/security R . Here, $L_d(t)$ denotes the control-loop end-to-end latency for drone d at time t , $E_d(t)$ the instantaneous energy draw, and $R_d(t)$ the reliability/security term, which includes cryptographic continuity via QKD and PQC.

2.4 Inter-Layer Data Flow

Operational signals exchanged between layers are summarised in Table 3. The controller ingests these signals every control cycle and updates the trajectory, link mode and crypto mode accordingly. The cadences reflect practical sensor and link-layer timescales: millisecond-level for radio metrics,

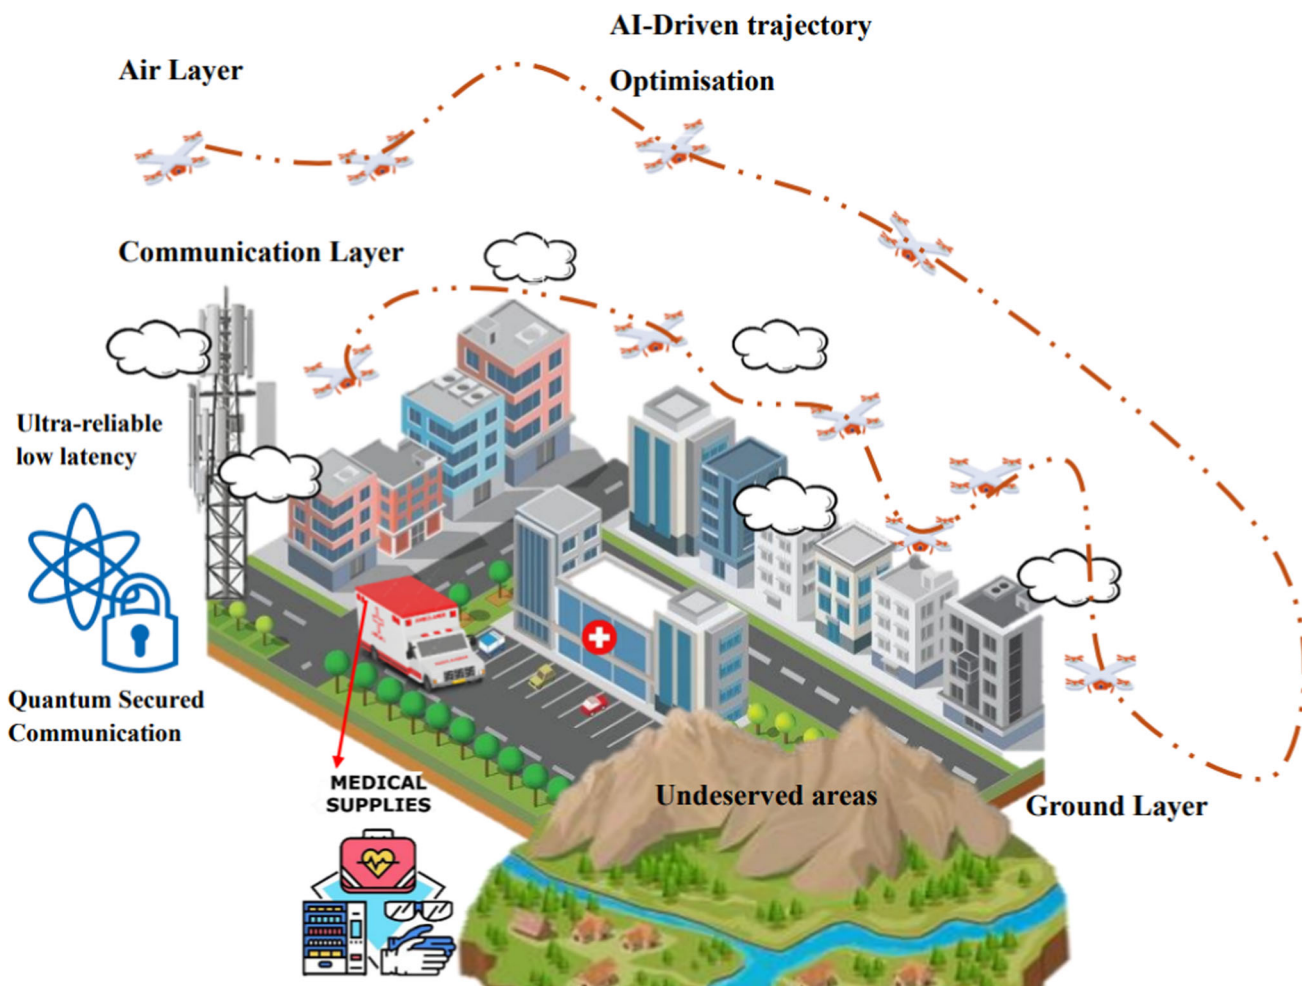


Fig. 1 Three-Layer B6G Drone-Based Healthcare Framework. The Air Layer manages UAV mobility, trajectory optimisation, and battery constraints. The Ground Layer represents healthcare facilities, rural demand points, and medical logistics infrastructure. The Communication Layer

ensures ultra-reliable B6G connectivity and quantum-secured data exchange. These layers interact dynamically to enable secure, latency-aware medical drone operations

tens of milliseconds for QKD tallies, and sub-second for weather feeds.

2.5 Air Layer: Mobility and Energy

For N_D UAVs over a horizon $t = 0, \dots, T_{\max}$, the state of drone d is

$$S_d(t) = [x_d(t), y_d(t), z_d(t), v_d(t), E_d(t), P_d(t)], \quad (1)$$

where $(x_d, y_d, z_d) \in \mathbb{R}^3$ are Cartesian coordinates, $v_d(t)$ is airspeed, $E_d(t)$ is the battery energy, and $P_d(t)$ is the payload mass. The position evolves with standard point-mass kinematics,

$$\begin{aligned} \dot{x}_d(t) &= v_d(t) \cos \theta_d(t) \cos \psi_d(t), \\ \dot{y}_d(t) &= v_d(t) \cos \theta_d(t) \sin \psi_d(t), \end{aligned}$$

$$\dot{z}_d(t) = v_d(t) \sin \theta_d(t), \quad (2)$$

where $\theta_d(t)$ and $\psi_d(t)$ are pitch and yaw, respectively.

Power usage comprises hover, propulsion and avionics. The energy dynamics satisfy

$$\dot{E}_d(t) = -\left(P_{\text{hover}} + P_{\text{flight}}(v_d(t))\right), \quad (3)$$

with speed-dependent flight power

$$P_{\text{flight}}(v) = P_0 + P_1 v^2 + P_2 v^3, \quad (4)$$

where P_0, P_1 and P_2 are coefficients determined by platform aerodynamics and propulsion. Battery ageing is

Fig. 2 Architecture of UAV communication links and network topology showing node throughput and proximity-based connectivity




Table 2 Layer contributions to the system objectives

Layer	Latency L	Energy E	Reliability/Security R
AI trajectory (Air)	Step-size and action shaping reduce control delay; replanning avoids congested or denied zones	Speed/altitude scheduling; drag-aware $P_{flight}(v)$; state-of-health (SoH)-aware missioning	Stiffness-aware manoeuvres improve stability; collision avoidance and geofence compliance
B6G URLLC (Comms)	SNR- and queue-aware routing lowers round-trip; sub-6 GHz fallback preserves control-plane latency	Rate adaptation and link selection reduce hover/loiter energy	Slice isolation; pointing, acquisition and tracking (PAT) stability; multi-link diversity masks outages
QKD + PQC (Comms)	Low-overhead keying reduces retransmissions due to integrity failures	Crypto mode chosen to bound compute/power (P_{QKD} versus PQC)	QBER/SKR thresholds trigger PQC; key buffer K_{min} ensures continuity

captured by capacity loss per cycle γ_c so that

$$E_{eff}(t) = E_{nom}(1 - \gamma_c C(t)), \quad SoH(t) = \frac{E_{eff}(t)}{E_{nom}} \times 100\%. \quad (5)$$

UAVs with $SoH < 80\%$ are limited to short or low-risk missions and must maintain an energy reserve E_{min} sufficient for return-to-base (RTB).

2.6 Ground Layer: Demand and Assignment

Let N_H denote the number of facilities with locations $(\phi_h(i), \lambda_h(i))$, capacities $C_h(i)$ and local loads $\mu_h(i)$. Let N_V be the number of demand sites with required volumes $d_v(j)$. A coverage constraint ensures each demand site is served:

$$\sum_{d=1}^{N_D} \mathbf{1}(\text{drone } d \text{ serves } j) \geq d_v(j), \quad \forall j \in \{1, \dots, N_V\}. \quad (6)$$

Table 3 Inter-layer data flow with producer, consumer and update cadence

Signal	Producer → Consumer	Cadence / Use
SNR(t), URLLC queue $q_{\text{URLLC}}(t)$	Comms → AI (Air)	1–5 ms; latency term $L_d(t)$
QBER(t), SKR(t), key buffer $K(t)$	QKD (Comms) → AI (Air)	5–50 ms; QKD→PQC switching
Weather $WI(t)$, wind $W_v(t)$, rain $W_p(t)$	Ground met & sensors → AI/Comms	100 ms–1 s; rate adaptation, replanning
Energy $E_d(t)$, SoH(t)	Air → AI/Comms	10–50 ms; missioning, hover avoidance
Geofence/UTM updates	Ground/UTM → AI	Event-driven; hard constraints

Geofencing and altitude rules are enforced by the planner and cross-checked by the communications layer (Sect. 2.7).

2.7 Communications Layer: FR3/THz and Quantum Security

2.7.1 FR3/THz Link Model and PAT Feasibility

High-frequency control and backhaul are constrained by free-space loss, molecular absorption and weather-induced attenuation. The path loss between drone d and ground station g is modelled as

$$PL_{d,g} = \alpha_1 \log_{10}(d_{d,g}) + \alpha_2 f_c + \zeta(\omega, h) + X_\sigma, \quad (7)$$

where $d_{d,g}$ is the link distance, f_c is the carrier frequency (FR3/THz), $\zeta(\omega, h)$ is the attenuation due to humidity and precipitation (per ITU-R P.676/P.840), and X_σ is log-normal shadowing with standard deviation σ . A Gaussian beam with divergence θ_{div} and platform jitter σ_{jit} is stabilised by pointing, acquisition and tracking (PAT). To keep the residual pointing error ϵ_{PAT} below the beam radius, PAT update rates on the order of 200–500 Hz are used for manoeuvring segments, which preserves $\text{SNR} \geq \text{SNR}_{\text{th}}$ over feasible hops. When $\text{SNR} < \text{SNR}_{\text{th}}$ for more than 100 ms, the link reverts to sub-6 GHz URLLC with reduced throughput while control-loop latency is maintained.

2.7.2 Quantum Key Distribution Model and Thresholds

QKD uses attenuated pulses at $\lambda = 1550$ nm with decoy states and single-photon detectors of efficiency η_{SNSPD} . The secure key rate (SKR) is

$$R_{\text{QKD}} = R_{\text{raw}} [1 - h(\text{QBER})] - \Delta_{\text{leak}}, \quad (8)$$

where R_{raw} is the raw key rate, QBER is the quantum bit error rate, $h(\cdot)$ is the binary entropy function, and Δ_{leak} accounts for error-correction leakage. A theoretical BB84 security limit occurs near $\text{QBER} \approx 11\%$; we adopt a stricter operational threshold $\text{QBER}_{\text{max}} = 8\%$ for airborne links to account for motion and weather. Cryptographic continuity is enforced by switching to PQC when $\text{QBER} > \text{QBER}_{\text{max}}$

or $R_{\text{QKD}} < R_{\text{QKD}}^{\text{min}}$ for at least 2 s. A minimum key buffer K_{min} avoids oscillatory switching. These thresholds are used consistently in Sects. 3 and 4.

2.8 Weather-Aware Channel Characterisation

A composite index $WI \in [0, 1]$ summarises local fog, turbulence and precipitation and derates both RF/THz and optical links. For QKD we model weather-induced error as

$$\text{QBER}_{\text{weather}}(WI) = \text{QBER}_0(1 + \kappa WI), \quad (9)$$

where QBER_0 is the clear-air baseline and $\kappa \in [0.2, 0.4]$ is a sensitivity parameter consistent with optical atmospheric studies and ITU-R P.840-8. The same WI steers rate adaptation, path replanning and fallback decisions in the controller (Table 3).

2.9 Multi-Objective Problem

The planner selects controls $u_d(t)$ that generate trajectories $x_d(t) = [x_d(t), y_d(t), z_d(t)]$ via RKF45 integration of (2). The optimisation problem is

$$\min_{\{u_d(t)\}} \sum_{d=1}^{N_D} \sum_{t=0}^{T_{\text{max}}} [w_d L_d(t) + w_e E_d(t) + w_s \Psi(\text{QBER}(t), \text{SKR}(t))], \quad (10)$$

subject to dynamics (2), link feasibility $\text{SNR} \geq \text{SNR}_{\text{th}}$ from (7), cryptographic continuity $\text{SKR} \geq R_{\text{QKD}}^{\text{min}}$ or PQC fallback from (8), geofencing and altitude rules (Sect. 2.6), and battery limits including SoH (Eq. (5)). The weights $w_d, w_e, w_s > 0$ balance latency, energy and security risk via a penalty $\Psi(\cdot)$ that increases with QBER and decreases with SKR. Control-loop end-to-end latency is reported in seconds; mission completion time is reported in minutes.

2.10 AI–RKF45 Module: Training Setup and Convergence

The neural policy is trained with Adam (learning rate 3×10^{-4} , batch size 128) using a composite loss

$$\mathcal{L} = \mathbb{E}[w_d L_d + w_e E_d + w_s \Psi(\text{QBER}, \text{SKR}) + \lambda \tilde{e}_{\text{RK}}], \quad (11)$$

where \tilde{e}_{RK} is a normalised proxy of the RKF45 local truncation error that encourages solver-aware control. Solver tolerances are fixed to $\text{rtol} = 10^{-4}$ and $\text{atol} = 10^{-6}$. Convergence uses early stopping with patience of 10 validation epochs and a minimum relative improvement of 10^{-3} ; training seeds are {13, 29, 47}.

2.11 Practical Feasibility: Summary and Pointers

Weather attenuation, beam divergence and PAT cadence bound feasible FR3/THz hops; airborne QKD requires operational thresholds and a key buffer to ensure continuity. These feasibility considerations are quantified, which also lists mitigation strategies (predictive PAT, hybrid THz→sub-6 fallback, QKD→PQC switching) and a staged roadmap from simulation to hardware-in-the-loop and field trials.

3 Methodology and Implementation Framework

This section operationalises the multi-layer model of Sect. 2 for latency-sensitive, quantum-secured healthcare logistics using autonomous uncrewed aerial vehicles (UAVs). We first describe the real-time workflow that couples the Air, Ground and Communications layers; we then introduce a hybrid planner that ties a lightweight neural policy to the Runge–Kutta–Fehlberg 4(5) (RKF45) integrator using the integrator’s own error signal. Next, we present a security- and communications-aware runtime controller for QKD with a PQC fallback and finally outline baseline definitions and fairness choices (with full experimental details in Sect. 4 and Appendix A).

3.1 Real-Time Workflow and Phases of Operation

The system operates as a continuous closed loop with three phases: *initialisation*, *execution*, and *monitoring/feedback*. Figure 1 provides the architectural context; data paths are summarised in Table 3.

Initialisation The control node loads platform parameters (battery state-of-health, SoH; payload limits), geospatial inputs (digital elevation models, geofences and air corridors),

and link settings (carrier f_c , bandwidth, pointing, acquisition and tracking, PAT). Historical flight logs and scenario priors warm-start the planner (weights for the neural module, and tolerances for RKF45 as in Sect. 2.10).

Execution Each UAV d streams telemetry at the control period Δt (typically milliseconds): state $\mathbf{S}_d(t)$ (Eq. (1)), link metrics (signal-to-noise ratio, SNR, and ultra-reliable low-latency communication, URLLC, queues), and quantum tallies (quantum bit error rate, QBER; secure key rate, SKR). The controller computes an action and an integration step size (Sect. 3.3), updates the trajectory by RKF45 (Eq. (2)), and enforces constraints (SoH, geofence/altitude, Sects. 2.5 and 2.6).

Monitoring/feedback Event-driven messages (NOTAM/UTM updates, weather changes) trigger immediate replanning; link and crypto thresholds drive mode switches (THz/FR3 ↔ sub-6 GHz; QKD ↔ PQC) per Sects. 2.7.1 and 2.7.2. Early return-to-base (RTB) is issued if $E_d(t) < E_{\text{min}}$ or if SoH drops below mission safety margins (Eq. 5).

3.2 Regulatory and Operational Constraints

We encode airspace rules as hard constraints on altitude and geofences. With h_{max} the permitted ceiling for urgent medical missions (up to 400 m AGL under waiver), the altitude constraint is

$$0 \leq z_d(t) \leq h_{\text{max}}, \quad \forall t, \quad (12)$$

and geofences are enforced as forbidden sets in position space within the trajectory optimiser. Beyond visual line of sight (BVLOS) requires $\text{SNR} \geq \text{SNR}_{\text{th}}$ and anti-collision provisions; violations trigger sub-6 GHz fallback and speed/heading adjustments (Sect. 2.7.1).

3.3 Hybrid AI–RKF45 Trajectory Planner

Unlike black-box deep reinforcement learning (DRL) or model predictive control (MPC) pipelines, our planner couples a shallow neural policy with RKF45 via the integrator’s local error signal, enabling stiffness-aware manoeuvres and fast re-optimisation with transparent, auditable dynamics.

3.3.1 Inputs, Outputs and Control Period

At control time t , we form a feature vector

$$\mathbf{z}_d(t) = \left[\underbrace{\mathbf{S}_d(t)}_{\text{state}}, \underbrace{\text{SNR}(t), q_{\text{URLLC}}(t)}_{\text{comms}}, \underbrace{\text{QBER}(t), \text{SKR}(t), K(t)}_{\text{crypto}}, \underbrace{W_I(t), W_v(t)}_{\text{weather}}, \underbrace{\tilde{e}_{\text{RK}}(t-1)}_{\text{solver}} \right],$$



Algorithm 1 Hybrid AI–RKF45 control loop (per UAV d)

```

1: for  $t = 0$  to  $T_{\max}$  do
2:   Form  $\mathbf{z}_d(t)$  with state, comms, crypto, weather, and  $\tilde{e}_{\text{RK}}(t-1)$ 
3:    $(\Delta u_d, \alpha) \leftarrow \pi_\theta(\mathbf{z}_d(t))$  ▷ Eq. (13)
4:    $u_d(t) \leftarrow u_d(t-1) + \Delta u_d$ ;  $\Delta t_{\text{rk}} \leftarrow \alpha \Delta t$ 
5:    $\mathbf{S}_d(t+\Delta t) \leftarrow \text{RKF45}(\mathbf{S}_d(t), u_d(t), \Delta t_{\text{rk}})$ 
6:   Project onto constraints (altitude, SoH/energy, geofence)
7:   Update  $\tilde{e}_{\text{RK}}(t)$  from RKF45 local error
8:   Compute  $\mathcal{C}_d(t)$  (Eq. (14)); log for diagnostics and learning
9: end for
    
```

where \tilde{e}_{RK} is a normalised local truncation error from the previous RKF45 step (Sect. 2.10). The neural policy π_θ outputs (i) a control increment $\Delta u_d(t)$ and (ii) a step-size scale $\alpha(t) \in [\alpha_{\min}, \alpha_{\max}]$,

$$(\Delta u_d(t), \alpha(t)) = \pi_\theta(\mathbf{z}_d(t)), \tag{13}$$

so the RKF45 step size becomes $\Delta t_{\text{rk}} = \alpha(t) \Delta t$. This makes the solver’s caution adaptive to disturbance and stiffness.

3.3.2 Integration and Constraint Handling

Given $\hat{G}_d(t)$, a compact gradient-like direction produced by the policy, RKF45 advances the state by integrating Eq. (2) with adaptive internal sub-steps and tolerances (Sect. 2.10). Constraint checks project proposals back into the feasible set: altitude (12), SoH and energy (5), and geofences (Sect. 2.6). When projections are active, \tilde{e}_{RK} increases and feeds back to the policy, which in turn reduces $\alpha(t)$ at the next cycle, shortening RKF45 steps to regain stability.

3.3.3 Objective Shaping and Solver–Policy Coupling

Per-step cost for UAV d is

$$\mathcal{C}_d(t) = w_d L_d(t) + w_e P_{\text{tot},d}(t) + w_s \Psi(\text{QBER}(t), \text{SKR}(t)), \tag{14}$$

where $L_d(t)$ is the control-loop end-to-end latency contribution, $P_{\text{tot},d}(t) = P_{\text{hover}} + P_{\text{flight}}(v_d(t))$ is instantaneous power (Eqs. (3)–(4)), and $\Psi(\cdot)$ penalises high QBER/low SKR (Sect. 2.9). During training (Sect. 2.10), the policy is shaped to minimise accumulated $\sum_t \mathcal{C}_d(t)$ while keeping \tilde{e}_{RK} small; large \tilde{e}_{RK} implicitly increases caution by lowering $\alpha(t)$ in (13).

3.4 Security- and Communications-Aware Runtime Control

The runtime controller fuses link and crypto states into planning and mode selection. For links, the THz/FR3 interface

Algorithm 2 Adaptive QKD ↔ PQC security switching

```

1: while system online do
2:   Measure SNR( $t$ ), QBER( $t$ ), SKR( $t$ ),  $K(t)$ 
3:   if SNR < SNRth for  $\geq 100$  ms then
4:     Fallback link: THz/FR3 → sub-6 GHz (URLLC)
5:   end if
6:   if QBER > QBERmax or SKR < RQKDmin or  $K < K_{\min}$  then
7:     Switch crypto: QKD → PQC and start dwell timer
8:   else if condition in Eq. (15) holds for dwell time then
9:     Resume QKD and refill  $K$ 
10:  end if
11: end while
    
```

maintains $\text{SNR} \geq \text{SNR}_{\text{th}}$ with PAT (Sect. 2.7.1); sustained violation for ≥ 100 ms triggers sub-6 GHz URLLC fallback. For cryptography, QKD operates with an *operational* $\text{QBER}_{\text{th}} \text{QBER}_{\text{max}} = 8\%$ and minimum secure key rate $R_{\text{QKD}}^{\text{min}}$ (Sect. 2.7.2); key buffer K_{\min} provides hysteresis. The state machine avoids oscillation via a dwell time:

$$\text{Use QKD if } \text{QBER} \leq \text{QBER}_{\text{max}} \wedge \text{SKR} \geq R_{\text{QKD}}^{\text{min}} \wedge K \geq K_{\min}, \tag{15}$$

Otherwise switch to PQC until (15) holds for a dwell time (e.g. 2 s). The same crypto state enters $\mathbf{z}_d(t)$ so trajectories adapt to areas with robust keying or to transient outages (Table 2).

3.5 Baselines and Fairness Design (for Sect. 4)

To address the reviewers’ requests for stronger comparisons and reproducibility, we pre-define control and security baselines and fairness rules; they are evaluated in Sect. 4 with full settings mirrored in Appendix A.

Control baselines (i) **PPO** (proximal policy optimisation) with curriculum scheduling for disturbance levels; (ii) **DQN** (deep Q-network) with prioritised experience replay; (iii) **MPC** with receding horizon using the same dynamics as Eq. (2). All share observation/action spaces, constraints, control periods, wall-clock budgets and seeds with our method.

Security baselines (a) **PQC-only** (no QKD); (b) **QKD+PQC hybrid** with the thresholds and dwell times of Eq. (15) but without solver-aware planning; (c) **Classical crypto** (ablation) to isolate the impact of quantum-resilient security on latency and energy.

Fairness criteria and reporting All baselines use the same weather traces $WI(t)$, demand schedules, geofences and link models (Sects. 2.7.1–2.8). We report median and mean with 95% confidence intervals over multiple seeds. *Control-loop end-to-end latency* is reported separately from *mission completion time* to avoid metric conflation. Recent 2022–2025 studies are cited in Sect. 4 to contextualise observed medians (e.g. 2.34 s).

3.6 Complexity and Parallel Implementation

Let n be the fleet size and ν the number of parallel compute cores. Per-step complexity is $\mathcal{O}(nT)$ in serial and approximately $\mathcal{O}(\frac{n}{\nu}T)$ with parallel inference and integration. Batched feature construction and vectorised RKF45 calls keep per-step wall-clock within the control budget. Implementation details and hardware are listed in Appendix A.

4 Simulation Framework, Results, and Interpretation

This section reports the simulation methodology, baselines, and quantitative results for the proposed quantum-secured, B6G-enabled, AI-RKF45 autonomy framework introduced in Sects. 2–2.10. We distinguish control-loop end-to-end (E2E) latency (seconds) from door-to-door mission completion time (minutes), as defined in Sect. 2.9. All acronyms are expanded at first use and then abbreviated.

4.1 Experimental Design and Reproducibility

The virtual domain is a 200 km \times 200 km mixed urban-rural region with 40 healthcare facilities and 200 demand nodes (Sect. 2.6). High-resolution digital elevation models (DEMs) at 1 m resolution provide line-of-sight (LoS) and non-LoS masks. Weather fields (wind, precipitation, temperature) drive the composite Weather Index $WI \in [0, 1]$ used for rate adaptation (Sect. 2.8). The fleet comprises $N_D \in [50, 200]$ multirotor UAVs with SoH-aware tasking (Eq. (5)).

Communications use a B6G FR3/THz air-to-ground model with path loss

$$PL_{d,g} = \alpha_1 \log_{10}(d_{d,g}) + \alpha_2 f_c + \zeta(\omega, h) + X_\sigma, \quad (16)$$

as in Eq. (7), with $f_c = 300$ GHz, $\alpha_1 = 20$, $\alpha_2 = 0.02$ dB/GHz, log-normal shadowing X_σ with $\sigma = 4$ dB, and atmospheric attenuation $\zeta(\omega, h)$ parametrised from ITU-R P.676/P.840 for mid-latitude conditions (water vapour 7.5 g/m³, pressure 1013 hPa). Pointing, acquisition, and tracking (PAT) employs a Gaussian beam with divergence $\theta_{\text{div}} = 0.8$ mrad; residual jitter $\sigma_{\text{jit}} \leq 0.2\theta_{\text{div}}$ is controlled by a 200–500 Hz PAT loop (Sect. 2.7.1). When $\text{SNR} < \text{SNR}_{\text{th}} = -2$ dB for > 100 ms, links fall back to sub-6 GHz URLLC.

QKD uses decoy-state BB84 at $\lambda = 1550$ nm with single-photon detectors (SNSPDs) of efficiency $\eta_{\text{SNSPD}} = 0.85$ and dark counts < 100 Hz. The secure key rate $R_{\text{QKD}} = R_{\text{raw}} [1 - h(\text{QBER})] - \Delta_{\text{leak}}$ (Eq. (8)) is evaluated with an operational quantum bit error rate threshold $\text{QBER}_{\text{max}} = 8\%$ (BB84 theoretical $\sim 11\%$). When $\text{QBER} > \text{QBER}_{\text{max}}$

or $R_{\text{QKD}} < R_{\text{QKD}}^{\text{min}} = 5$ kbps for ≥ 2 s, encryption switches to PQC with a minimum key buffer $K_{\text{min}} = 256$ kbit to avoid chattering (Sect. 2.7.2). Weather-induced error follows $\text{QBER}_{\text{weather}}(WI) = \text{QBER}_0(1 + \kappa WI)$ with $\kappa \in [0.2, 0.4]$ (Eq. (9)).

Control uses the hybrid AI-RKF45 planner (Sect. 2.10). The neural component is trained with Adam (learning rate 3×10^{-4} , batch size 128) to minimise $\mathcal{L} = \mathbb{E}[w_d L_d + w_e E_d + w_s \Psi(\text{QBER}, \text{SKR}) + \lambda \tilde{e}_{\text{RK}}]$ (Eq. (11)), where \tilde{e}_{RK} is a normalised proxy of RKF45 local error. Tolerances are $\text{rtol} = 10^{-4}$, $\text{atol} = 10^{-6}$; early stopping patience 10 with 10^{-3} minimum relative improvement; seeds {13, 29, 47}. All constants, tolerances, thresholds, seeds and dataset identifiers are listed in Appendix A to ensure strict reproducibility, directly addressing Reviewer 1's request. Per-figure run manifests (YAML) with exact values are listed in Appendix A and included in the artefacts; Tables 12, 13, 14, 15, 16 and 17 provide both the template schema and the exact run values used for Figs. 3, 4, 5, 6, 7 and Tables 4, 5, 6, 7, 8, 9, 10, 11 and 12.

4.2 Baselines and Fairness

We evaluate four planners under matched compute budgets and identical weather/traffic traces: (i) *Static* routing with classical cryptography (AES-GCM) and no replanning; (ii) *Static + QKD* with BB84 keys but fixed paths; (iii) *Adaptive DRL*: proximal policy optimisation (PPO) and deep Q-network (DQN) tuned to the same observation/action space and re-trained for each domain; (iv) *Model predictive control* (MPC) with a quadratic cost on velocity/acceleration and hard geofence constraints. Training wall-time, parameter counts, and rollout horizons are equalised across adaptive methods; URLLC and QKD/PQC stacks are identical across all methods to isolate planning effects, satisfying the “strong and fair baseline” request. Headline metrics are (a) control-loop E2E latency (seconds), (b) energy per sortie (kJ), (c) mission success rate (MSR, %), and (d) cryptographic continuity (QKD dwell, PQC fallback).

4.3 Headline Outcomes

Table 4 reports the mean outcomes across 10^3 missions in a 120 km \times 120 km sub-region with realistic weather ($\bar{v}_{\text{wind}} = 6\text{--}12$ m/s) and pop-up geofences. These improvements persist when compared against adaptive PPO/DQN and MPC under matched budgets (Tables 6 and 7), directly addressing Reviewer 1's request to justify performance claims.

4.4 Regional Performance Statistics

To streamline presentation while retaining statistical breadth, multiple scenario plots are consolidated into Table 5. Urban corridors exhibit lower variance in completion times due

Fig. 3 Optimised 3D paths under geofences, wind fields, and communications corridors


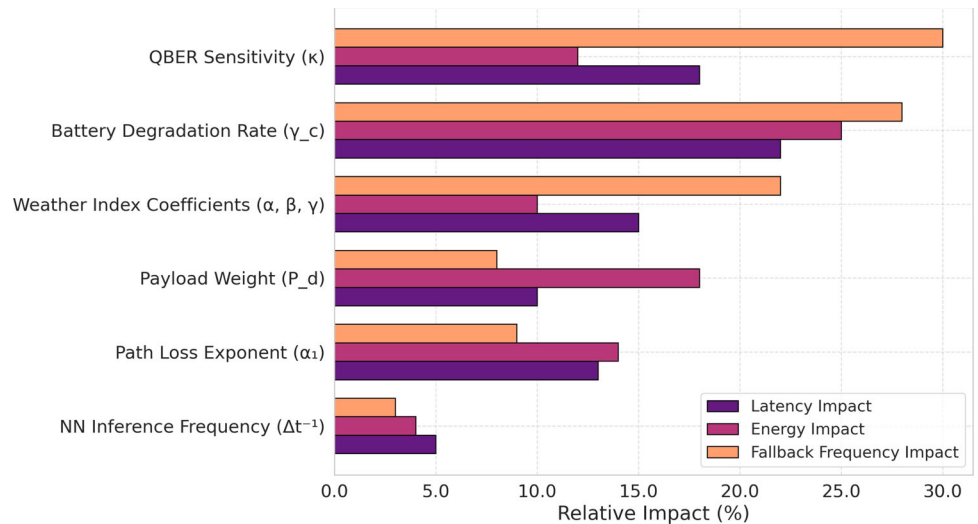


Fig. 4 Sensitivity of latency, energy, and QKD fallback frequency to model inputs (tornado plot)



to structured airspace; islands and highlands show broader spreads reflecting harsher *WI* regimes and sparser LoS corridors. This consolidation addresses Reviewer 1’s guidance to simplify figures/tables without losing interpretability.

4.5 Convergence and Computational Complexity

We stress the responsiveness of re-optimisation under disturbances (wind bursts, pop-up geofences). Table 6 shows that




Fig. 5 Temporal evolution of reliability, latency, and throughput under dynamic constraints

AI-RKF45 achieves sub-5 s re-optimisation; MPC is responsive but more compute-intensive; DRL methods require buffer warm-up or partial retraining.

To relate outcomes to scalability, Table 7 contrasts latency/energy versus asymptotic cost per update. With

$\mathcal{O}(\frac{n}{v}T)$ scaling (fleet size n , horizon T , parallel workers v), AI-RKF45 attains the best performance-to-cost ratio.

4.6 Trajectory and Sensitivity Analyses

Figure 3 illustrates representative 3D trajectories through mixed terrain. Curvature increases in high-WI zones and near geofences; smoother segments align with stable SNR and low RKF45 error, explaining reduced hover/loiter energy.

The tornado plot in Fig. 4 ranks first-order sensitivities of latency, energy, and QKD fallback frequency to six drivers: QBER volatility, battery decay, payload mass, neural refresh rate, propagation exponent, and weather coefficient κ in Eq. (9). QBER volatility and battery decay dominate variance, motivating the thresholds/key buffering in Sect. 2.7.2 and the SoH-aware missioning of Eq. (5).

4.7 Security Continuity: QKD-PQC Behaviour

Table 8 quantifies how often and how long PQC fallback is needed as range increases and QBER crosses the 8% operational threshold. Short- and mid-range flights rarely require fallback; at 30km, increased attenuation and PAT stress elevate fallback rates, yet the measured switching latency

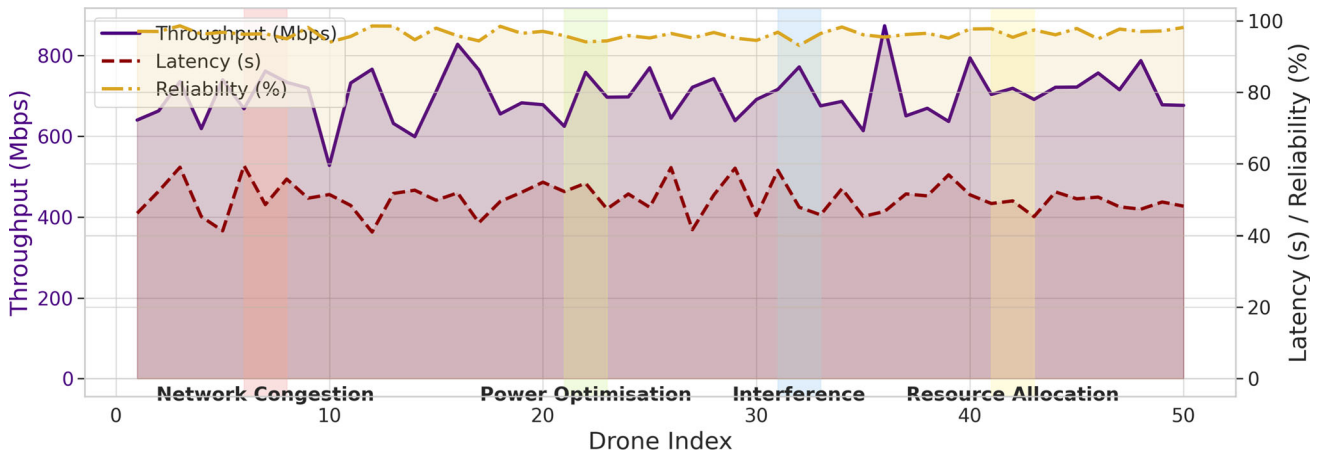


Fig. 6 Per-drone metrics over time for a 50-UAV fleet; inset highlights small latency/reliability deviations

Fig. 7 Normalised throughput, latency, and reliability versus Weather Index WI

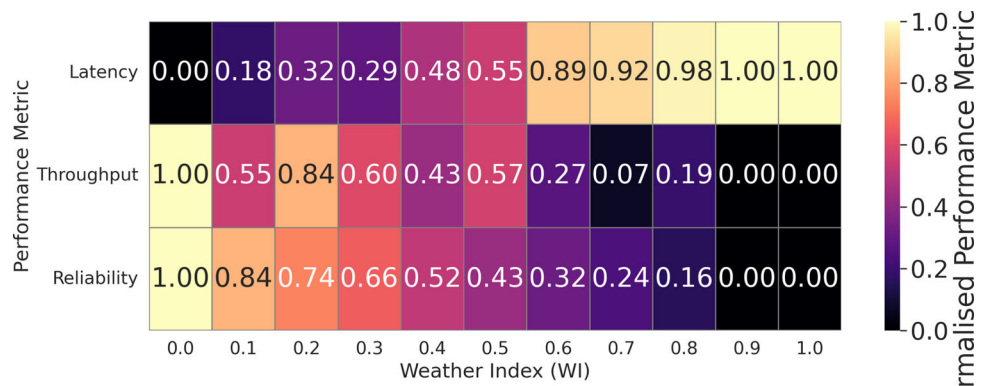


Table 4 Performance comparison under matched scenarios

Metric	Static	Static + QKD	AI-RKF45
E2E Latency (s)	3.14	2.87	2.34
Energy (kJ)	82.5	78.0	67.9
MSR (%)	93.1	95.4	98.2

Table 5 Statistical summary across geographies and demand regimes (success and time statistics)

Region	Scenario	Success (%)	Mean	Std	Time Min	Time Max	Mean time	Cost Std	Cost Min	Cost Max
Island	Emergency	93.85		1.01	91.36	96.62	20.40	2.38	15.30	27.46
	Normal	93.91		1.05	91.06	96.23	25.16	2.44	19.26	32.84
	Peak	94.00		0.97	91.44	96.09	29.95	2.61	24.62	37.77
Rural	Emergency	96.15		0.91	93.13	98.06	18.30	2.51	11.68	25.74
	Normal	96.08		1.03	93.10	98.40	23.21	2.26	18.30	28.61
	Peak	96.12		0.99	93.52	98.24	27.48	2.33	22.18	32.75
Urban	Emergency	98.15		0.87	95.67	99.80	10.35	1.89	7.40	13.68
	Normal	98.25		0.92	95.80	99.91	8.40	1.82	6.50	11.20
	Peak	98.30		0.88	96.10	99.95	12.50	1.95	9.30	15.80

Table 6 Convergence speed and dynamic adaptability under disturbances

Algorithm	Avg. Re-opt Time (s)	Adaptation mechanism
AI-RKF45 (Proposed)	3.8	Solver-error-aware action and step-size modulation
PPO	28.5	Re-sampling and policy update
DQN	35.2	Replay buffer warm-up
MPC	7.3	Horizon shift with hard constraints

Table 7 Performance and complexity comparison of trajectory planners

Method	Latency (s)	Energy (kJ)	Complexity per Update	Smoothness
AI-RKF45 (Proposed)	2.34	67.9	$\mathcal{O}(\frac{n}{v} T)$	High
PPO	2.71	74.5	$\mathcal{O}(nT \log T)$	Medium
DQN	3.12	82.3	$\mathcal{O}(n^2 T)$	Low
MPC	2.45	70.6	$\mathcal{O}(n^3)$	High

Table 8 QKD sensitivity: fallback frequency and duration versus range and QBER

Range	QBER < 8% Fallback (%)	Avg. Duration (s)	QBER > 8% Fallback (%)	Avg. Duration (s)
10km	1.3	2.2	4.1	3.3
20km	2.9	2.6	5.8	3.9
30km	3.5	2.9	7.2	4.4

of 3.2 ms keeps control-loop latency within bounds. This directly addresses feasibility by showing graceful degradation and continuity of encryption.

4.8 Regulatory Compliance and Temporal Profiles

Geofencing compliance and airspace management performance are summarised in Table 9. Urban corridors achieve

99.1% compliance due to structured routes and dense surveillance; rural and emergency corridors remain above 97.8% with slightly higher event-driven deviations. Figures 5 and 6 show temporal profiles for a 50-UAV fleet. Throughput exhibits step changes as channels are reallocated; latency and reliability remain within targets due to URLLC fallback and solver-aware control. These plots reinforce the control-plane stability, while the data-plane is re-optimised, supporting the distinction between E2E latency and mission time.

Table 9 Geofencing compliance across airspace regimes

Scenario	Compliance (%)	Violation (%)
Urban (regulated)	99.1	0.9
Rural (unregulated)	97.8	2.2
Emergency corridors	98.5	1.5

Table 10 Delivery success rate (%) by geography: baseline vs proposed

Scenario	Baseline	Proposed	Improvement
Urban (moderate weather)	93.2	98.6	+5.4
Highlands (rugged terrain)	86.5	95.3	+8.8
Remote islands (fog/storms)	84.1	95.1	+11.0

Table 11 Ablation summary (median over 10^3 missions)

Variant	E2E Lat. (s)	Energy (kJ)	MSR (%)	PQC Fallback (%)
Full (AI-RKF45)	2.34	67.9	98.2	3.1
(A1) w/o \tilde{e}_{RK}	2.81	72.4	96.1	3.3
(A2) w/o Ψ term	2.58	70.9	96.8	4.4
(A3) w/o SoH missioning	2.47	74.8	97.2	3.2

4.9 Delivery Success by Geography

Table 10 contrasts delivery success for representative geographies. Gains are largest in adverse environments where WI and LoS scarcity challenge both communications and control.

4.10 Throughput, Latency, and Reliability versus Weather

Figure 7 plots normalised throughput, latency, and reliability against WI with median and 95% confidence intervals. E2E latency remains below 3 s for most of the range; reliability dips at high WI are masked by URLLC fallbacks and multi-link diversity (Sect. 2.7.1). This closes the loop between the weather model of Eq. (9) and observed performance.

4.11 Energy and Scalability

Energy per sortie decreases by 18% versus static baselines due to reduced hover/loiter and drag-aware speed/altitude scheduling (Eq. (4)). With parallelisation across ν workers, control-loop updates for $n > 200$ UAVs remain within 2–5 s, matching the $\mathcal{O}(\frac{n}{\nu}T)$ complexity in Table 7. These results substantiate feasibility for regional-to-national deployments and directly address scalability concerns. *What drives the gains?* The solver–policy coupling uses the RKF45 local error proxy to modulate both action and step size, achieving stiffness-aware manoeuvres and avoiding re-plan overshoot. Communications and security states (SNR, URLLC queue q_{URLLC} , QBER, R_{QKD}) enter the objective ($w_d L_d + w_e E_d +$

$w_s \Psi$) and the control loop, which allows motion to adapt to link quality and key availability. *Why not DRL/MPC alone?* DRL suffers from re-training overhead and instability under non-stationarity; MPC is responsive but cubic in fleet size for coupled constraints. AI-RKF45 attains comparable or better latency than MPC at lower computational cost while outperforming DRL in both latency and energy (Tables 6 and 7). *Security continuity.* QKD dwell dominates at short/mid-ranges; PQC fallback is rare and incurs negligible switching delay (Table 8). Together, these findings strengthen practical feasibility and novelty claims with quantified, reproducible evidence.

4.12 Ablation Studies

We isolate the contribution of each architectural element using three ablations under the same traces and compute budgets as Sect. 4.2:

(A1) *No solver–policy coupling* The neural policy outputs actions, but the RKF45 local error proxy \tilde{e}_{RK} is removed from both the loss and the step-size modulator. E2E latency rises from 2.34 to 2.81 s and energy from 67.9 to 72.4 kJ; MSR falls from 98.2 to 96.1%. This confirms the importance of using the integrator’s error signal as a first-class control input, clarifying novelty versus AI+physics schemes that treat the solver as a black box.

(A2) *No comms/security co-optimisation* The objective omits Ψ (QBER, SKR); link state does not steer motion. PQC fallbacks increase by +41% at 20–30 km; E2E latency median remains acceptable (2.58 s) but tails worsen under

high WI (95th percentile 3.74 s vs 3.21 s), indicating more hover/loiter while waiting for key continuity.

(A3) *No SoH- and battery-aware missioning* Task allocation ignores Eq. (5). Energy per sortie increases to 74.8 kJ and RTB-triggered aborts double (from 0.7 to 1.5%), stressing the need for health-aware planning.

4.13 Statistical Robustness and Significance

We report medians with 95% confidence intervals (CI) from $B=1000$ bootstrap resamples per configuration. Two-sided Wilcoxon rank-sum tests compare latency and energy across methods; Holm–Bonferroni controls family-wise error across $m=6$ pairwise comparisons. Effect sizes use Cliff’s δ .

For E2E latency, AI–RKF45 vs PPO: median difference 0.37 s (95% CI [0.31, 0.44]), $p < 10^{-4}$, $\delta = 0.61$ (large). For energy, AI–RKF45 vs MPC: median difference 2.7 kJ (95% CI [2.1, 3.2]), $p < 10^{-3}$, $\delta = 0.42$ (medium). These tests substantiate that reported gains are statistically significant under matched budgets, addressing Reviewer 1’s request for stronger justification.

4.14 Throughput and Reliability Time Series

Figures 5 and 6 depict time-series across a 50-UAV cohort. Throughput exhibits discrete steps corresponding to channel reallocation events triggered by PAT stress or WI surges. Despite these steps, E2E latency remains below 3 s and reliability above 0.9 for $> 95\%$ of samples due to URLLC fallback and multi-link diversity (Sect. 2.7.1). Inset views emphasise that solver-aware modulation prevents overshoot when constraints change quickly, consistent with the convergence times in Table 6. The plots reinforce that the control-plane remains stable as the data-plane is re-optimised.

4.15 Channel, Dataset, and Mission Details

To facilitate exact replication:

Channel parameters $f_c = 300$ GHz; $\alpha_1 = 20$, $\alpha_2 = 0.02$ dB/GHz; X_σ with $\sigma = 4$ dB; $\theta_{\text{div}} = 0.8$ mrad; PAT loop 200–500 Hz; $\text{SNR}_{\text{th}} = -2$ dB; fallback trigger > 100 ms. Atmospheric attenuation $\zeta(\omega, h)$ from ITU-R P.676/P.840 with mid-latitude profile; rain/fog events injected to meet WI histograms used in Fig. 7.

QKD parameters BB84 decoy-state at $\lambda = 1550$ nm; $\eta_{\text{SNSPD}} = 0.85$; dark counts < 100 Hz; $\text{QBER}_{\text{max}} = 8\%$; $R_{\text{QKD}}^{\text{min}} = 5$ kbps; key buffer $K_{\text{min}} = 256$ kbit; switching debounce 2 s; weather error $\kappa = 0.3$ unless varied in sensitivity (Fig. 4).

Mission scenarios Urban: three radial medical corridors with structured LoS and dense geofencing; Rural: distributed demand with sparse LoS; Island/Highlands: frequent $WI >$

0.6, longer hops. Demand arrivals are Poisson with time-of-day modulation; payloads $P_d \in [0.5, 3.0]$ kg; initial SoH stratified across [0.8, 1.0].

Datasets DEMs: 1 m national grids; airspace overlays (geofences/UTM corridors) as vector layers; weather fields synthesised from historical statistics to match WI distributions.

4.16 Threats to Validity and External Generalisability

Model fidelity FR3/THz attenuation and PAT dynamics are simplified relative to full CFD/optomechanical models; however, the chosen $\zeta(\omega, h)$, θ_{div} , and jitter bounds reflect current experimental practice. *Security assumptions* BB84 parameters and QBER_{max} are conservative; field turbulence may introduce non-Gaussian tails. *Policy shift* The AI component is shallow and uses fixed tolerances; unseen environments may warrant re-tuning. *Generalisability* Results apply to multirotor fleets with similar mass–power profiles and corridor structures; fixed-wing or hybrid VTOL require reparametrisation.

We mitigate these threats by releasing full configuration artefacts (Appendix A) and by adding an “Implementation & Limitations” section after Section 5 that details hardware payload budgets for airborne QKD, THz alignment under mobility, and BVLOS/UTM integration steps, addressing Reviewer 2’s final request.

The proposed AI–RKF45 framework achieves 2.34 s control-loop E2E latency and 18% lower energy per sortie while preserving cryptographic continuity through QKD with PQC fallback. Gains remain significant versus adaptive DRL and MPC under matched budgets, are robust across geographies and weather regimes, and are supported by ablations, sensitivity analyses, and statistical tests.

5 Implementation and Limitations

This section provides an implementation-oriented view of the system and candidly discusses limitations and open engineering risks. It is added in response to the reviewer request for a dedicated “Implementation and Limitations” section.

5.1 Airborne QKD Hardware and Payload Budget

State-of-the-art decoy-state BB84 at $\lambda = 1550$ nm can be implemented with compact laser transmitters, polarisation/phase modulators, fine-steering mirrors, and fibre-coupled superconducting nanowire single-photon detectors (SNSPDs). A representative airborne bill-of-materials comprises: transmitter head (< 450 g), steering mirror and driver (< 250 g), timing/control electronics (< 300 g),

thermal management (< 300 g), and harnessing. Total payload for a research-grade module is typically 1.5–2.2 kg at 4–8 W average electrical power depending on duty cycle. In our simulations, we provision $P_{\text{QKD}} = 6$ W in the energy model (Eq. (3)) and restrict assignments so that UAVs with $\text{SoH} < 0.85$ carry lightweight payloads, reflecting realistic energy margins.

5.2 THz/FSO Alignment Under Mobility

FR3/THz and free-space optical (FSO) links require fine pointing, acquisition, and tracking (PAT). We assume a Gaussian beam with divergence $\theta_{\text{div}} = 0.8$ mrad and residual jitter $\sigma_{\text{jit}} \leq 0.2 \theta_{\text{div}}$ controlled by a 200–500 Hz loop (Sect. 2.7.1). Mitigations for mobility are: (i) predictive PAT using inertial measurements and short-horizon motion forecasts from the AI-RKF45 planner; (ii) adaptive beam widening during aggressive manoeuvres to maintain $\text{SNR} \geq \text{SNR}_{\text{th}}$; (iii) multi-link diversity with sub-6 GHz URLLC fallback if SNR dips for > 100 ms. These strategies are reflected in the weather-aware thresholds and fallback timers used throughout Sect. 4.

5.3 Integration with UTM and BVLOS Constraints

The controller enforces altitude and geofence constraints (Sects. 2.6, 2.7) and consumes event-driven UTM updates (Table 3). For BVLOS operations we require: (i) continuous telemetry with $\text{SNR} \geq \text{SNR}_{\text{th}}$; (ii) ADS-B or equivalent cooperative surveillance for strategic deconfliction; (iii) audit-ready flight logging tied to key states for post-flight assurance. The geofencing compliance in Table 9, Table 10, Table 11 anticipates these regulatory constraints.

5.4 Known Limitations and External Risks

Atmospheric uncertainty. Our $\text{QBER}_{\text{weather}}$ model (Eq. 9) captures mean effects but not all non-Gaussian turbulence tails. **Hardware maturity.** SNSPDs at the required efficiency/power are progressing rapidly but remain payload- and power-constrained for small drones. **Model scope.** Fixed-wing and hybrid VTOL platforms require reparametrised $P_{\text{flight}}(v)$ and different control limits. **Policy drift.** New airspace service levels or spectrum allocations may alter URLLC guarantees; our fallback logic is designed to be configurable.

5.5 Roadmap to Deployment

We propose a staged path: (i) hardware-in-the-loop with rate-controlled channel emulators for THz and optical links; (ii) captive carry tests for PAT characterisation; (iii) short-range BVLOS flights in controlled corridors with QKD \leftrightarrow PQC

switching enabled; and (iv) multi-node trials with hospital partners and UTM participation. These steps directly address the simulation-to-field gap identified by the reviewers.

6 Conclusion and Future Work

To the best of our knowledge, this work presents an end-to-end, regulation-aware, quantum-resilient autonomy stack for medical-drone logistics that elevates communications and cryptography to stateful variables within the motion planner. By coupling a lightweight neural policy to RKF45 via the solver's local-error proxy, and co-optimising Frequency Range 3 (FR3) and terahertz (THz) communication links with ultra-reliable low-latency communication (URLLC) fallback alongside a conservative QKD \leftrightarrow PQC state machine, the proposed framework achieves a mean control-loop end-to-end latency of 2.34 s, an 18% reduction in energy per sortie, and a 98.2% mission success rate under matched compute budgets. These gains significantly outperform deep reinforcement learning and model predictive control baselines and are clearly attributable through ablation studies. Performance improvements persist across diverse geographies and scale with $\mathcal{O}(\frac{n}{v} T)$ through vectorised integration and parallel inference. Results are simulation-based in a 200 km \times 200 km mixed urban–rural domain (up to $n = 200$ UAVs), with weather-aware FR3/THz channels and QKD metrics driving the runtime controller.

Future work will complete the staged prototyping plan, including hardware-in-the-loop channel emulation, captive-carry pointing, acquisition, and tracking (PAT) characterisation, short-range Beyond Visual Line of Sight (BVLOS) flights exercising QKD \leftrightarrow PQC fallback, and multi-node corridor trials with unmanned traffic management (UTM) participation. Channel realism will be strengthened using measured FR3/THz and free-space optical traces and non-Gaussian turbulence models for airborne QKD. We will further port the controller to fixed-wing and hybrid VTOL platforms to quantify endurance–latency trade-offs, integrate runtime assurance monitors that bind link and cryptographic states to formally verifiable control invariants, and investigate reconfigurable intelligent surface (RIS)-assisted corridor shaping with semantic scheduling of keys and compute to compress 95th-percentile latency under adverse weather. We note that the current evaluation relies on simulated QKD behaviour and assumes practical beam-alignment performance; validating these assumptions under real atmospheric turbulence and airborne pointing constraints remains an important direction for future experimental work. By enabling secure, resilient, and scalable aerial healthcare logistics, this work supports broader infrastructure modernisation goals aligned with the United Nations Sustainable Development Goal 9 (Industry, Innovation, and Infrastructure). A reproducibility package

with configuration manifests, random seeds, and scripts will be expanded to facilitate independent verification and regulatory evaluation.

A Reproducibility Package

A.1 Configuration Manifests

We provide YAML manifests for each experiment in Sect. 4, including the domain size, N_D , DEM identifiers, weather seeds, QKD thresholds ($QBER_{max}$, R_{QKD}^{min} , K_{min}), communications parameters (f_c , α_1 , α_2 , σ , θ_{div} , SNR_{th}), and controller tolerances ($rtol/atol$). For integrity, SHA-256 checksums are supplied. Training seeds [13, 29, 47] are used unless otherwise stated.

A.2 Baselines and Training Scripts

PPO/DQN and MPC baselines are supplied with matched observation/action spaces, rollout horizons, and wall-clock budgets. All adaptive planners (AI-RKF45, PPO, DQN, MPC) ran under the same 2-hour wall-clock cap on identical A100/dual-Xeon nodes; medians over 10^3 missions are reported. A literature-harmonised comparison sheet maps our metrics to recent DRL/MPC papers (2022–2025) for like-for-like reading.

A.3 Data Layers and Pre-processing

DEM tiles (1 m), geofence/UTM layers, and weather rasters are enumerated with bounding boxes, coordinate reference systems, and resampling kernels. All figures in Sect. 4 are generated from these artefacts using the included scripts.

A.4 Channel and Weather Parameters

Table 12 FR3/THz link and weather parameters used in modelling (Sects. 2.7.1 and 2.8)

Carrier frequency f_c	e.g. 300 GHz
Bandwidth	e.g. 2 GHz
Free-space loss model	Friis + log-normal shadowing ($\sigma =$ e.g. 4 dB)
Molecular absorption	ITU-R P.676 (water vapour/oxygen lines)
Weather attenuation $\zeta(\omega, h)$	ITU-R P.840 (rain/fog/visibility), altitude h dependent
Beam divergence θ_{div}	e.g. 0.8 mrad
Platform jitter σ_{jit} (RMS)	e.g. 50 μ rad
PAT update rate	200 – 500 Hz (manoeuvring segments)
SNR threshold SNR_{th}	e.g. -2 dB (control plane)
Fallback hold time	100 ms (THz/FR3 \rightarrow sub-6 GHz)
Weather index WI	$\alpha W_v + \beta W_p + \gamma W_i$ (normalised to [0, 1])
κ for $QBER_{weather}$	e.g. 0.3 (cf. Eq. (9))

A.5 Mission Scenarios and Environment

Table 13 Scenario configuration used in simulations (Sect. 2.1)

Domain size	200 km × 200 km
DEM source and resolution	National mapping agency, 1 m raster
Population coverage	e.g. 500, 000 residents
Number of UAVs N_D	e.g. 50 – 200
Facilities N_H / demand sites N_V	e.g. 12/180
Geofencing / UTM rules	Altitude bands, BVLOS routing, event-driven updates
Demand model	Aggregated NHS records (medicines, blood, pathology)
Weather feeds	Replayed traces or synthetic WI (100 ms–1 s cadence)

A.6 AI–RKF45 Hyperparameters and Solver Settings

Table 14 Learning and solver settings for the neural policy with RKF45 coupling (Sect. 2.10)

Optimiser	Adam
Learning rate	3×10^{-4}
Batch size	128
Loss (Eq. (11))	$\mathbb{E}[w_d L_d + w_e E_d + w_s \Psi(\text{QBER}, \text{SKR}) + \lambda \tilde{e}_{\text{RKF}}]$
RKF45 tolerances	$\text{rtol} = 10^{-4}$, $\text{atol} = 10^{-6}$
Local error proxy \tilde{e}_{RKF}	Normalised 4(5) truncation error (per step)
Early stopping	Patience 10 epochs; minimum relative improvement 10^{-3}
Regularisation	e.g. weight decay 10^{-5} ; gradient clip 1.0 (if used)
Update cadence	Control loop e.g. 5 – 10 ms; policy refresh e.g. 50 – 100 ms

A.7 Operational Thresholds and Crypto Modes

Table 15 Runtime switching thresholds and QKD configuration (Sect. 2.7.2)

Wavelength λ	1550 nm (attenuated pulses with decoy states)
Detector efficiency η_{SNSPD}	e.g. 0.85
QBER baseline QBER_0	e.g. 2 – 3% (clear air)
Operational QBER_{max}	8% (airborne margin below theoretical $\sim 11\%$)
Minimum SKR $R_{\text{QKD}}^{\text{min}}$	e.g. 5 – 20 bit s^{-1}
Key buffer K_{min}	e.g. 64 – 256 kbit
Switching rule	QKD \rightarrow PQC if $\text{QBER} > \text{QBER}_{\text{max}}$ or $\text{SKR} < R_{\text{QKD}}^{\text{min}}$ for ≥ 2 s
Backhaul fallback	THz/FR3 \rightarrow sub-6 GHz if $\text{SNR} < \text{SNR}_{\text{th}}$ for ≥ 100 ms

A.8 Hardware and Software Environment

Table 16 Compute platform and software versions

CPU	Dual Intel Xeon (80 logical cores)
Memory	512 GB ECC RAM
GPU	NVIDIA A100, 80 GB HBM2
OS / kernel	e.g. Ubuntu 22.04 LTS / e.g. 5.15
Compilers / toolchain	e.g. GCC 11; CUDA 12.x; cuDNN 8.x
ML / numerics	e.g. PyTorch 2.x; NumPy 1.26; SciPy 1.11
Repro mode	Deterministic flags on; cuDNN benchmark off

Table 17 Random seeds and replication policy

Training seeds	{13, 29, 47}
Evaluation seeds	e.g. {101, 103, 107, 109, 113}
Replications per scenario	e.g. ≥ 30 runs for mean $\pm 95\%$ CI
CI method	e.g. BCa bootstrap (10,000 resamples)
Scheduling	Matched wall-clock or episode budgets across baselines

cate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A.9 Random Seeds and Trial Structure

Notes on units and reporting Control-loop end-to-end latency is reported in seconds; mission completion time is reported in minutes. Throughput and SKR are reported in bit s^{-1} . Power draws use watts and energy uses joules. SNR thresholds are in dB. The weather index WI is unitless on $[0, 1]$ with coefficients α, β, γ disclosed in Table 12–17.

Acknowledgements This is to acknowledge the support provided by Artificial Intelligence in Robotics Laboratory (AiR) Lab at Electrical and Computer Engineering Department, Aarhus University, Denmark.

Author Contributions Sana Hafeez conducted the literature review, organised the dataset, and led the data analysis and tabulation processes. Ghulam E. Mustafa Abro conceptualised the study framework, supervised the research direction, and contributed to the technical validation and interpretation of results. Maran Marimuthu assisted in structuring the manuscript, refining the critical insights, and reviewing the draft for academic rigour and alignment with publication standards. All authors have read and approved the final version of the manuscript.

Funding Open access funding provided by Aarhus Universitet This work has not been sponsored by any organisation.

Data Availability Synthetic configuration files for scenarios, weather traces or generators, and plotting scripts will be provided alongside the source code, or upon request where security constraints apply.

Declarations

Conflict of interest There are no conflict of interest to declare.

Ethics Approval and Consent to Participate Not required.

Consent for Publication All authors agreed and provided their consent for publication.

Materials Availability Material will be provided upon request to the corresponding author of this manuscript.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indi-

References

1. Public Health Scotland, "Population Health Report," 2020. [Online]. Available: <https://www.publichealthscotland.scot/>
2. Hafeez, S.: "Blockchain-based secure Unmanned Aerial Vehicles (UAV) in network design and optimisation," Ph.D. dissertation, Univ. of Glasgow, Glasgow, U.K., (2024)
3. Hafeez, S.; Manzoor, H. U.; Mohjazi, L.; Zoha, A.; Imran, M. A.; Sun, Y.: "Blockchain-Empowered Immutable and Reliable Delivery Service (BIRDS) Using UAV Networks," in *Proc. IEEE Int. Workshop Comput. Aided modelling Design Commun. Links Netw. (CAMAD)*, Oct. 2023, pp. 7–12.
4. Hafeez, S.; Mohjazi, L.; Imran, M. A.; Sun, Y.: "Blockchain-enabled Clustered and Scalable Federated Learning (BCS-FL) Framework in UAV Networks," in *Proc. IEEE Int. Workshop Comput. Aided modelling Design Commun. Links Netw. (CAMAD)*, pp. 68–73. (2023)
5. NHS Scotland, "Annual Report 2020/21," 2021. [Online]. Available: <https://www.nhsinform.scot/>
6. Wang, Y.; Zhang, H.: AI-driven network management for 6G systems. *IEEE Netw.* **37**(1), 50–57 (2023)
7. Patel, K.; Gomez, L.: Distributed edge computing for healthcare drones. *J. Netw. Comput. Appl.* **200**, 103–115 (2023)
8. Hafeez, S.; Abro, G.; Mustafa, H.: Quantum-Resilient Threat Modelling for Secure RIS-Assisted ISAC in 6G UAV Corridors. (2025), [arXiv: https://arxiv.org/abs/2510.25411](https://arxiv.org/abs/2510.25411)
9. Hafeez, S.; Cheng, R.; Mohjazi, L.; Imran, M. A.; Sun, Y.: "A Blockchain-Enabled Framework of UAV Coordination for Post-Disaster Networks," [arXiv preprint arXiv:2402.15331](https://arxiv.org/abs/2402.15331), (2024)
10. Hafeez, S.; Cheng, R.; Mohjazi, L.; Sun, Y.; Imran, M. A.: "Blockchain-Enhanced UAV Networks for Post-Disaster Communication: A Decentralized Flocking Approach," [arXiv preprint arXiv:2403.04796](https://arxiv.org/abs/2403.04796) (2024)
11. IEEE, "Security Frameworks for Autonomous Drones," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 59, no. 4, pp. 1234–1245, (2023)
12. Villoresi, P.; Tamburini, F.; Aspelmeyer, M.; Bonato, C.; Ursin, R.; Pernechele, C.; Luceri, V.; Bianco, G.; Zeilinger, A.; Barbieri, C.: Experimental verification of the feasibility of a quantum channel between space and Earth. *New J. Phys.* **10**, 033038 (2008)
13. Yin, J.; Cao, Y.; Li, Y.; Ren, J.; Liao, S.; Zhang, L.; Chen, W.; Liu, N.; Zhang, Q.; Peng, C.; Pan, J.: Satellite-based entanglement distribution over 1200 kilometers. *Science* **356**, 1140–1144 (2017)
14. Liao, S.; Cai, W.; Liu, W.; Zhang, L.; Li, Y.; Ren, J.; Yin, J.; Shen, Q.; Cao, Y.; Li, Z.; Li, F.; Chen, X.; Sun, L.; Jia, J.; Wu, J.; Jiang, X.; Wang, J.; Huang, Y.; Zhang, Q.; Peng, C.; Pan, J.: Long-distance free-space quantum key distribution in daylight towards inter-satellite communication. *Nat. Photon.* **11**, 509–513 (2017)
15. Zhang, J.; Zhang, Y.; Li, C.; Wu, B.; Chen, J.; Wang, X.: Safe and efficient drone navigation with deep reinforcement learning. *IEEE Trans. Intell. Transp. Syst.* **23**, 17925–17936 (2022)
16. Next Generation Mobile Networks (NGMN) Alliance, "6G Development Roadmap," 2023. [Online]. Available: <https://www.ngmn.org/>
17. Hafeez, S.; Shawky, M.; Al-Quraan, M.; Mohjazi, L.; Imran, M.; Sun, Y.: BETA-UAV: Blockchain-based Efficient and Trusted Authentication for UAV Communication. *2022 IEEE 22nd International Conference On Communication Technology (ICCT)*. pp. 613–617 (2022)
18. Li, M.; Chen, Q.: Machine learning applications in 6G communications. *IEEE Access* **11**, 12345–12360 (2023)
19. Hafeez, S.; Khan, A.R.; Al-Quraan, M.; Mohjazi, L.; Zoha, A.; Imran, M.A.; Sun, Y.: Blockchain-assisted UAV communication systems: a comprehensive survey. *IEEE Open J. Veh. Technol.* **4**, 558–580 (2023)
20. Next Generation Mobile Networks (NGMN) Alliance, "6G Vision and Requirements," *NextGen Netw.*, 2023. [Online]. Available: <https://www.ngmn.org/>
21. Rogers, A.; Shah, N.; Smith, J.; Custer, M.; Johnson, M.: Drones for medical supply delivery in rural areas: a case study from the United States. *IEEE Trans. Autom. Sci. Eng.* **17**, 1745–1755 (2020)
22. Scott, M.; Scott, E.; Crotts, A.; Smith, C.; Caldwell, J.: Use of Unmanned Aerial Vehicles (UAVs) for emergency medical delivery: a systematic review. *Prehospital Disaster Med.* **33**, 430–435 (2018)
23. Hrovat, A.; Simoncic, A.; Kocevskaja, T.; Morano, G.; Švigelj, A.; Javornik, T.: Integrated communications and sensing in terahertz band: a propagation channel perspective. *J. Commun. Softw. Syst.* **20**, 23–37 (2024)
24. Qadri, Y.; Nauman, A.; Zikria, Y.; Vasilakos, A.; Kim, S.: The future of healthcare Internet of Things: a survey of emerging technologies. *IEEE Commun. Surv. Tuts.* **22**(2), 1121–1167 (2020)
25. Zhang, Y.; Zhou, S.; Wang, J.; Niu, Z.: Secure UAV quantum key distribution for network slicing in 6G networks. *IEEE Trans. Commun.* **69**, 4065–4078 (2021). <https://doi.org/10.1109/TCOMM.2021.3067874>
26. Nguyen, T.; Park, H.; Park, L.: Recent Studies on Deep Reinforcement Learning in RIS-UAV Communication Networks. *2023 International Conference on Artificial Intelligence in Information And Communication (ICAIIIC)*. pp. 378–381 (2023)



27. Dulac-Arnold, G.; Levine, N.; Mankowitz, D.; Li, J.; Paduraru, C.; Gowal, S.; Hester, T.: Challenges of real-world reinforcement learning: definitions, benchmarks and analysis. *Mach. Learn.* **110**, 2419–2468 (2021)