

Article

Hierarchical Multi-Agent Navigation Through the 72-h Thermal Drift Cliff

Mosab Alrashed , Humoud Aldaihani  and Mohammad Alqattan 

School of Engineering and Computing, American International University (AIU), Saad Al Abdullah, Jahra 91103, Kuwait; h.aldaihani@aiu.edu.kw (H.A.); m.alqattan@aiu.edu.kw (M.A.)

* Correspondence: m@dralrashed.com or m.alrashed@aiu.edu.kw

Highlights

What are the main findings?

- The analytical 72-h navigation cliff is reproduced in simulation to within 0.8% ($t_c = 72.9 \pm 2.4$ h across 2430 Monte Carlo runs) and explained by irreversible thermal hysteresis creep in the MEMS gyroscope. Under the same channel statistics, a CTMC-coupled collaborative SLAM filter holds GPS-denied localization within the 2 m budget, where Swarm-SLAM, Kimera-Multi, and D2SLAM all exceed 3.9 m.
- When paired with cognitive radio hopping, a 3D Gaussian process RF map queried inside the predictive controller raises link availability from 12.4% to 81.7% under barrage jamming and cuts RF map error by 48.6%. An AHP-weighted multi-objective cost reduces latency by 15.1% and jitter by 12.9% at negligible SNR cost, all inside the 100 ms Layer 2 Jetson Xavier NX budget.

What are the implications of the main findings?

- Communication quality, RF map structure, and channel stochasticity can be folded into the cost of predictive control as first-class planning variables, providing a tractable theory for long-endurance multi-UAV mission planning under intermittent and adversarial links. The resulting channel quality–fleet size envelope allows a planner to read off the minimum fleet that meets a target sortie.
- The simulation-validated bifurcation at $t_c = 72.9 \pm 2.4$ h, together with $t_c^{\text{eff}} > 96$ h for $N = 5$ at $\pi_{\text{good}} = 0.80$, supplies a quantitative design baseline for next-generation GPS-denied UAV systems. All claims are simulation-based; hardware-in-the-loop testing and long-duration field deployment remain priority future work.

Abstract

Long-endurance unmanned aerial vehicle (UAV) missions beyond 72 consecutive flight hours face a reliability boundary at which thermal gyroscope drift drives the inertial navigation system (INS) position error into rapid nonlinear divergence at a predictable threshold t_c . This paper presents BAZ II, a simulation-validated multi-agent navigation system that extends the analytical BAZ (bifurcation-aware zonal navigation) framework. Its central idea is to treat communication quality as a planning resource and combine it with multi-agent collaboration, making the navigation cliff a manageable degradation event rather than a hard operating limit. Four contributions support this idea: a thermal-hysteresis MEMS gyroscope drift model reproduces the analytical cliff in simulation and supplies its physical mechanism; a distributed collaborative simultaneous localization and mapping (SLAM) filter coupled to a stochastic continuous-time Markov chain (CTMC) inter-agent channel sustains GPS-denied localization within the operational accuracy budget; a 3D Gaussian process RF-aware model predictive controller (MPC) with cognitive radio



Academic Editor: Pablo Rodríguez-González

Received: 6 May 2026

Revised: 21 June 2026

Accepted: 26 June 2026

Published: 24 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

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

frequency-hopping restores link availability under jamming, while an analytic hierarchy process (AHP)-weighted multi-objective communication cost improves latency and jitter at negligible signal-to-noise ratio cost; finally, the integrated controller executes within the onboard real-time budget of an NVIDIA Jetson Xavier NX. All results are obtained in simulation, with hardware-in-the-loop and field testing remaining as priority future work.

Keywords: thermal gyroscope drift; INS bifurcation threshold; collaborative SLAM; Gaussian process radio mapping; cognitive radio frequency hopping; model predictive control; MEMS inertial navigation; multi-agent UAV resilience

1. Introduction

1.1. Background and Motivation

Long-duration unmanned aerial vehicles (UAVs)—defined here as platforms capable of multi-day sorties exceeding 48 consecutive flight hours—have become essential to autonomous environmental monitoring, persistent area surveillance, and high-altitude communication relay [1–4]. Unlike rotary-wing platforms where mission times are typically limited to 2–8 h, long-endurance UAVs must maintain navigation capability across multiple thermal cycles per day [5,6], tolerate extended Global Positioning System (GPS) outages, and operate over highly variable communication links—conditions that lie well outside the design envelope of conventional navigation and control architectures. Throughout this paper, the acronym BAZ (Bifurcation-Aware Zonal navigation) denotes the precursor analytical framework of Alrashed et al. [7], while BAZ II denotes the simulation-validated multi-agent extension developed here. The precursor study [7] introduced the idea that communication quality should be treated as a planning resource and proactively shaped through trajectory selection. The BAZ study established three theoretical results: (1) ϵ -optimality of a sigmoid penalty model predictive control (MPC) algorithm under continuous-time Markov chain (CTMC) communication uncertainty ($\epsilon \leq 0.12$); (2) exponential stability of a hysteresis-modulated switched architecture ($\lambda \geq 0.23$); and (3) analytical derivation of a 72-h bifurcation threshold. MPC is adopted as the unifying control framework here because it is the only standard receding-horizon method that natively combines the three properties required by long-endurance multi-UAV navigation: (i) hard pointwise constraints on attitude rate, thrust, and obstacle clearance; (ii) stochastic disturbance models such as the CTMC channel state, which enter as expectations over the prediction horizon; and (iii) position-dependent cost terms—specifically, the GP-queried SNR field J_{comm} and the SLAM information cost J_{info} in Equation (1)—that have no closed-form analogue in linear quadratic regulators or sliding-mode formulations. Therefore, the BAZ II contribution is not the MPC paradigm itself but the integration of these three terms into a single computationally tractable cost function on embedded hardware. Beyond the 72-h cliff, thermally induced gyroscopic drift drives the INS position error through a sharp transition from bounded linear growth to rapid nonlinear divergence. Accelerometer thermal sensitivity also exists [8], but is sub-dominant on the multi-day timescale: typical tactical-grade MEMS accelerometer thermal bias of $\sim 0.5 \text{ mg}/^\circ\text{C}$ produces a Schuler-bounded position error oscillation that does not accumulate, whereas the integrated gyroscope drift (Equation (5)) grows quadratically and is the rate-determining failure mode beyond $\sim 48 \text{ h}$. Therefore, the accelerometer error budget is folded into the INS noise model in Section 3.3 and does not require a separate bifurcation analysis. BAZ was validated through Monte Carlo simulation over 2430 missions consisting of 54,686 flight hours [7] and demonstrated a 243% endurance extension together with sub-9 m circular error probable (CEP) navigation precision even

under severe communication degradation ($q_{\text{comm}} < 0.3$). Two limitations of BAZ motivate the present study: first, the analytical 72-h cliff was not connected to a physically based simulation model, an omission flagged by the BAZ peer review process; second, BAZ did not consider multi-agent coordination or 3D spatially resolved RF planning. Both limitations are addressed here. A summary of the key BAZ results on which this work builds is provided in Table 1 (located in Section 3 for ease of reference).

Table 1. BAZ key results (Alrashed et al. [7]). ε : MPC sub-optimality; λ : Lyapunov stability rate; t_c : analytical navigation cliff threshold (this row reports the BAZ Phase 1 analytical value, 72.3 ± 1.8 h; the BAZ II simulation-validated value of 72.9 ± 2.4 h is reported separately in Equation (9) and in Section 3.3); CEP: circular error probable; CTMC: continuous-time Markov chain.

Result	Value	Significance
MPC ε -optimality	$\varepsilon \leq 0.12$	Near-optimal under ergodic CTMC
Lyapunov stability rate	$\lambda \geq 0.23$	Exponential stability confirmed
Analytical cliff threshold	$t_c = 72.3 \pm 1.8$ h	Derived from MEMS drift model
Navigation precision (degraded)	CEP < 9 m	Sub-9 m even at $q_{\text{comm}} < 0.3$
Endurance extension	243% over baseline	Validated over 2430 MC missions

1.2. Research Gap and Contributions

While the BAZ analytical framework [7] establishes the existence and theoretical character of the 72-h navigation cliff, several gaps remain before the framework can support reliable multi-agent deployments at long endurance. A first-principles physical mechanism for the cliff has not been demonstrated; the analytical CTMC model has not been coupled to a multi-agent SLAM front-end; the scalar reactive communication penalty offers no spatial foresight in 3D urban environments; and the original communication cost functional treats SNR as the sole proxy for link quality, ignoring bandwidth, latency, and jitter. These four gaps motivate the contributions of the present paper.

The central research question addressed here is the following: by treating communication quality as a first-class planning resource within a cooperative multi-agent system, can the analytically predicted 72-h navigation cliff be converted from a hard endurance limit into a managed degradation event that a mission planner can anticipate, and can this be achieved within the real-time computational budget of embedded UAV hardware? The main scientific contribution of the paper is a simulation-based affirmative answer to this question: a single hierarchical navigation system, BAZ II, in which a physically grounded cliff model, collaborative SLAM, RF-aware predictive control, and multi-objective communication cost together extend the effective cliff time while remaining inside the onboard real-time limit. The four primary contributions listed below are the components that, in combination, answer this question.

In response to second-round reviewer feedback that an earlier draft presented seven items at uneven levels of novelty, the contribution list has been consolidated into four primary contributions, supplemented by two implementation enablers that are explicitly not claimed as new (Sections 3.5 and 4.1.5). The four primary contributions are summarized below; the prose is deliberately non-parallel, since the contributions differ in character (a validated physical mechanism, an estimator, an integrated controller, and a cost-functional reformulation) rather than being four instances of a single template.

C1. Physical mechanism for the 72-h navigation cliff. The BAZ analytical bifurcation [7] predicted, but did not physically explain, the abrupt growth in INS position error beyond ~ 72 h. A thermal hysteresis MEMS gyroscope drift model is constructed here from first principles—irreversible micromechanical creep accumulated over successive diurnal cycles—and is calibrated against published tactical-grade datasheets [9–11]. The bifurcation mechanism of Equation (7) reproduces the analytical cliff to within 0.8% over

2430 Monte Carlo runs; this is the first time an analytical 72-h cliff is shown to coincide with a physically derivable failure mode rather than a property of the analytical model alone.

C2. Stochastic communication-aware collaborative SLAM. The single-agent BAZ formulation cannot fuse landmark observations across multiple agents and offers no mechanism for the cliff to be arrested by external position aiding. Section 3.4 introduces an information-weighted multi-agent SLAM filter that is coupled directly to the inter-agent channel through a two-state CTMC, lumped from the BAZ three-state model. Crucially, the ϵ -optimality and Lyapunov stability guarantees inherited from BAZ are preserved under the lumped channel. On a common 30-min visual-inertial benchmark against Swarm-SLAM, Kimera-Multi, and D2SLAM (Section 3.4), the proposed filter is the only one for which the RMS localization error remains within the 2 m operational budget under stochastic CTMC degradation.

C3. 3D RF-aware predictive control with cognitive radio resilience. The scalar reactive penalty of BAZ steers the vehicle only after an SNR threshold has been crossed. In its place, Sections 3.5 and 3.6 integrate a 3D Gaussian process (GP) RF map built using an anisotropic spatiotemporal kernel and queried directly inside the MPC cost J_{comm} with a cognitive radio (CR) frequency-hopping layer that operates against barrage and sweep jammers. The combination is what is novel here: GP-based predictive routing alone cannot survive adversarial jamming, and CR alone cannot exploit spatial RF structure. Their integration restores link availability from 12.4% to 81.7% under barrage jamming and lowers RF map prediction error by 48.6% relative to an isotropic baseline.

C4. Multi-objective communication cost via the analytic hierarchy process. The original BAZ communication cost is scalar in SNR. Section 4.1.3 replaces it with an analytic hierarchy process (AHP) scalarization that admits bandwidth, latency, and jitter as additional axes; the pairwise-comparison matrix is provided in Supplementary Material S1. The reformulated cost is shown via tornado analysis to be insensitive to $\pm 20\%$ weight perturbations, dispelling the concern that the AHP reduction is a fragile choice of weights, and yields a 15.1% latency reduction and 12.9% jitter reduction with negligible SNR loss.

The two implementation enablers are deliberately separated from the contribution list, since neither is novel. Online recursive GP updates (Section 3.5, $m = 64$ inducing points at 10 Hz) follow the sparse-GP machinery of [12]; the parameter choices are documented in Supplementary S5. The recursive update is required for tractable execution of C3 on the Jetson Xavier NX, but the algorithmic structure itself is not new. A structure-exploiting condensed Riccati QP step (Section 4.1.5) is the textbook block-tridiagonal factorization [13] that also underlies HPIPM [14] and FORCES [15]; authorship of the recursion itself is not claimed. What is reported is the integration of this factorization as the real-time iteration step of the BAZ II MPC, and the resulting solve-time profile that keeps the fully integrated controller inside the 100 ms Layer 2 budget on embedded hardware.

Table 2 maps each primary contribution to the specific prior works that it extends, enabling readers from each sub-community to locate the most relevant advances. All four primary contributions are unified within a computationally constrained hierarchical architecture (Section 4.1) that manages the tradeoff between processing fidelity and onboard computational resources. The adaptive sparse/full mode switching qualifies this as a multi-mode navigation system [13], dynamically balancing high-fidelity and resource-conserving operation based on mission criticality. The switching architecture relies on stability theory for switched and hybrid systems [16,17] and nonlinear control fundamentals [18], with rigid-body attitude tracking [19] providing the geometric consistency framework.

Table 2. Map of contributions to the literature for the four primary BAZ II contributions (C1–C4) and the two implementation enablers (E1–E2). Each row identifies the prior work that the contribution or enabler directly extends. GP: Gaussian process; SLAM: simultaneous localization and mapping; MPC: model predictive control; CR: cognitive radio; CTMC: continuous-time Markov chain.

BAZ II Item	Directly Extends	Key References
Primary contributions		
C1. Thermal-hysteresis cliff validation	BAZ CTMC bifurcation analysis	[7,9,11]
C2. Distributed collaborative SLAM under stochastic comm.	D ² SLAM, Swarm-SLAM, Kimera-Multi	[20–23]
C3. 3D RF-aware predictive MPC + cognitive-radio resilience	SNR-constrained planning, GP radio mapping, dynamic spectrum access	[24–29]
C4. AHP-weighted multi-objective communication cost	Scalar SNR penalty (BAZ)	[7,30,31]
Implementation enablers (not claimed as new contributions)		
E1. Online recursive sparse-GP updates	Sparse-GP machinery	[12,32]
E2. Structure-exploiting condensed Riccati QP step	HPIPM, FORCES, textbook Riccati factorization	[13–15]

2. Related Work

2.1. INS Thermal Drift and Multi-UAV SLAM

Tactical-grade MEMS gyroscopes are dominated by thermal sensitivity, with position error growing quadratically under bias instability [9,33,34]. Temperature drift compensation using machine learning methods has been explored for MEMS accelerometers [8,35] and inertial navigation systems [36], while attitude and quaternion representations for strapdown INS are well-established [37,38]. GPS-denied navigation remains a fundamental challenge for long-endurance UAVs, addressed through vision-based approaches [39–45], lightweight simultaneous localization and mapping (SLAM) for resource-constrained platforms [46], celestial navigation alternatives [47], autonomous landing [48], and Lie group-based visual localization [49]. The EuRoC benchmark dataset [50] and Stanford drone dataset [51] have become standard evaluation testbeds, while Cadena et al. [52] provide a comprehensive overview of SLAM past, present, and future.

Modern UAV navigation estimators rely on extensions of the Kalman filter, primarily the extended Kalman filter (EKF) and unscented Kalman filter (UKF), which are the appropriate references for the nonlinear strapdown navigation problem considered in this paper [52,53]; filter divergence remains a concern in long-horizon operation [53]. GNSS reliability under degraded conditions has been characterized via Monte Carlo methods [54]. While multi-UAV SLAM systems like Kimera-Multi, Swarm-SLAM, and decentralized pose graph optimization [21,22,55] have matured, they do not explicitly address sustaining collaboration during the multi-day nonlinear error regime predicted in BAZ [7]. Vision- and language navigation surveys [56] and autonomous landing on maritime vessels [57] extend the operational envelope further.

2.2. Communication-Aware Path Planning

RF awareness in UAV planning has been explored via SNR-constrained optimization and radio map-based signal-to-interference-plus-noise ratio (SINR) planning [24,26,58]. UAV communication for wireless networks has been studied extensively [59–61], including energy minimization for rotary-wing UAVs [62], joint communication and trajectory optimization for IoV applications [63], and time-variant multipath channel models for UAV-to-ground links [64]. Finite-state Markov channel models underpin the CTMC framework used in this work [65]. Communication-aware swarm surveillance has been studied for hierarchical multi-UAV architectures [66]; adaptive control via communication information optimization [67] and distributed MPC for channel resource allocation [68] represent closely related contributions. Asynchronous distributed trajectory optimization for UAV swarms

under communication constraints has also been proposed [69]. GP regression is a standard tool for spatial RF mapping [25], but achieving 10 Hz updates on embedded hardware while managing the cyber–physical tradeoffs of long-endurance missions remains an open challenge that is addressed in this work.

GP-guided MPC trajectory planning with these techniques for UAVs remains largely unexplored. This work brings together GP-based 3D RF mapping, cognitive radio (CR) anti-jamming, and position-dependent CTMC channel models in a single MPC framework, thereby shifting communication management from reactive to predictive and directly extending the CTMC model developed in BAZ [7].

3. Materials and Methods

3.1. System Architecture

BAZ II extends the three-layer hierarchical architecture of BAZ [7] with three new modules: a collaborative SLAM front-end, a 3D RF-aware predictive controller, and a 72-h cliff monitor (Figure 1). The reference platform is a 23.8 kg hexarotor equipped with tactical-grade MEMS sensors and partitioned across three operating layers: Layer 1 (Navigation, 100 Hz), Layer 2 (Planning, 10 Hz), and Layer 3 (Mission Control, <1 Hz). Airframe attitude and rate dynamics follow standard rigid-body formulations [70,71], with geometric tracking control [72,73] providing attitude consistency across the operating modes. The handling quality envelope conforms to MIL-STD-1797A [74,75]; the IMU performance envelope is benchmarked against navigation-grade SBG Systems units [76], and the platform mass and payload allocation are sized against the published DJI Matrice 400 specification [77]. Detailed specifications and temporal scales are provided in Section S5.3.1 of Supplementary Material S5.

3.2. Unified MPC Cost Function

The integrated MPC minimizes

$$J = \underbrace{\sum_{k=0}^{N_p} \|\mathbf{x}_k - \mathbf{x}_k^{\text{ref}}\|_Q^2}_{\text{tracking}} + \underbrace{\sum_{k=0}^{N_p-1} \|\mathbf{u}_k\|_R^2}_{\text{control}} + \underbrace{J_{\text{comm}}(\mathbf{x}_{0:N_p})}_{\text{communication}} + \underbrace{J_{\text{info}}(\boldsymbol{\Sigma})}_{\text{information gain}} + \underbrace{J_{\text{comp}}(\mathbf{c})}_{\text{compute budget}} \quad (1)$$

where $N_p = 20$ is the prediction horizon (BAZ used $N_p = 50$ at 10 Hz with a 5-s lookahead [7]; the reduced horizon reflects the higher per-step cost of 3D GP map queries and SLAM covariance propagation on the NVIDIA Jetson Xavier NX, consistent with the SWaP-constrained limitation identified in BAZ). $\mathbf{Q}$ and $\mathbf{R}$ are state and control weighting matrices inherited from BAZ [7], and the three new cost terms are as follows:

$$J_{\text{comm}}(\mathbf{x}_{0:N_p}) = w_c \sum_{k=0}^{N_p} \max(0, \gamma_{\text{SNR}} - \text{SNR}(\mathbf{x}_k))^2 \quad (2)$$

$$J_{\text{info}}(\boldsymbol{\Sigma}) = w_i \sum_{k=0}^{N_p} \text{tr}(\boldsymbol{\Sigma}_k^{\text{SLAM}}) \quad (3)$$

$$J_{\text{comp}}(\mathbf{c}) = w_\kappa \sum_{k=0}^{N_p} \max(0, c_k - c_{\text{max}})^2 \quad (4)$$

where $\text{SNR}(\mathbf{x}_k)$ is queried from the 3D RF map, $\boldsymbol{\Sigma}_k^{\text{SLAM}}$ is the SLAM covariance, and c_k denotes the normalized computational load at prediction step k . The weights are set to

$w_c = 1.0$, $w_i = 0.5$, $w_\kappa = 2.0$, and $\gamma_{\text{SNR}} = 10$ dB, determined via pairwise comparison following the analytic hierarchy process [31]. The pairwise comparison matrix (consistency ratio $\text{CR} < 0.10$) is provided as Supplementary Material S1. Because the sub-costs have heterogeneous physical dimensions (m^2 , dB^2 , dimensionless), each weight w absorbs the necessary unit conversion so that J is dimensionless; therefore, reported MPC cost values are relative to the baseline configuration. The communication cost J_{comm} is a direct spatial extension of the scalar sigmoid penalty from BAZ [7], preserving the ε -optimality guarantee of Theorem 1 in that work.

3.3. Simulation Validation of the 72-h Navigation Cliff

3.3.1. Thermal Hysteresis Gyroscope Drift Model

The gyroscope drift rate is modeled as the superposition of four components:

$$\dot{\theta}(t) = b_0 + b_{\text{rw}}(t) + b_{\text{th}}(T(t)) + b_{\text{hyst}}(t) \quad (5)$$

where:

- $b_0 \sim \mathcal{N}(\mu_{\text{BI}}, \sigma_{\text{BI}}^2)$ is the turn-on bias with $\mu_{\text{BI}} = 1.0$ deg/h and $\sigma_{\text{BI}} = 0.15$ deg/h (tactical-grade MEMS specification [9,10,78]);
- $b_{\text{rw}}(t) = \int_0^t w(\tau) d\tau$ is the gyro random walk, $w(t) \sim \mathcal{N}(0, \sigma_{\text{rw}}^2)$ with $\sigma_{\text{rw}} = 0.002$ deg/h^{1/2};
- $b_{\text{th}}(T) = \kappa_T \cdot (T(t) - T_0)$ is the linear thermal sensitivity with $\kappa_T = 0.008$ deg/h/°C;
- $b_{\text{hyst}}(t) = \kappa_H \int_0^t \left| \frac{dT}{d\tau} \right| d\tau$ is the thermal hysteresis accumulation term with $\kappa_H = 1.8 \times 10^{-4}$ deg/h per cumulative °C h⁻¹, representing irreversible micromechanical creep in the MEMS sensing element [9,11,79–82].

The temperature profile models desert diurnal cycling:

$$T(t) = T_0 + T_A \sin\left(\frac{2\pi t}{T_P} - \frac{\pi}{2}\right) + \eta(t) \quad (6)$$

with $T_0 = 35^\circ\text{C}$, $T_A = 15^\circ\text{C}$, $T_P = 24$ h, and $\eta(t)$ a filtered Gaussian noise process ($\sigma_\eta = 1.0^\circ\text{C}$, 1 h correlation time chosen to match the characteristic timescale of mesoscale atmospheric fluctuations [2]). This sinusoidal model has a long history in atmospheric boundary layer research [2,83], achieving less than 2°C RMS approximation error relative to standard measured diurnal temperature profiles for arid environments [83,84]. Figure 2 illustrates the evolution of the four drift components over a 120-h simulation, showing the progressive accumulation of b_{hyst} across successive diurnal cycles and its nonlinear interaction with the total drift rate.

3.3.2. Bifurcation Mechanism

The key observation is that the hysteresis term $b_{\text{hyst}}(t)$ grows monotonically with each diurnal cycle—it captures irreversible micromechanical creep in the MEMS sensing element caused by repeated thermal stress. After approximately three complete diurnal cycles (≈ 72 h), b_{hyst} exceeds a bifurcation threshold β^* , causing the effective drift rate to transition from its nominal value to a rapidly amplified regime




Figure 1. Three-layer hierarchical system architecture of BAZ II, rotated to a full-page landscape layout to enlarge the module labels. The white boxes with solid borders are BAZ modules that were inherited from [7]; the teal-green dashed boxes represent new BAZ II modules introduced in this paper. Layer 1 (100 Hz) processes sensor data to perform collaborative SLAM and sensor fusion; Layer 2 (10 Hz) plans MPC trajectories that take into account communication constraints; Layer 3 (<1 Hz) is used for mission management and for adapting the amount of available computational resources.




Figure 2. Validation of the simulation for the 72-h navigation cliff ($n = 2430$ Monte Carlo runs). (a) Decomposition of Gyro Bias Error from a single representative run: constant offset b_0 , random walk b_{rw} , thermal b_{th} , hysteresis b_{hyst} , and total bias showing increasing hysteresis creep as function of time. (b) Position error envelope, linear scale (0–96 h): P5–P95 band (orange), interquartile range (IQR) P25–P75 (blue filled region), median P50 (blue); the gray dashed line represents linear extrapolation to 72 h, the shaded area is the cliff transition zone (68–78 h), and values are shown at 48/72/96 h. (c) Same position error envelope but on log scale (0–120 h): the mean $\bar{e}$ (green dashed line) illustrates superlinear post-cliff error increase. (d) Histogram of cliff onset times t_c with Gaussian fit $\mathcal{N}(72.9, 2.4^2)$ h and kernel density estimate (KDE); $\pm 1\sigma$ band shaded.

$$\dot{e}_{\text{pos}}(t) = |\dot{\theta}(t)| \cdot \left(1 + \alpha \cdot \max(0, b_{\text{hyst}}(t) - \beta^*)^2\right), \quad (7)$$

where $\alpha = 8000$ is the amplification coefficient encoding the unit conversion between gyroscope drift rate (deg/h) and position error rate (km/h); see Supplementary S5, Section S5.1 for the calibration procedure. Here, $\beta^* = 0.0213 \pm 0.0006$ is the per-unit bifurcation threshold, with inter-unit manufacturing variability drawn from $\mathcal{N}(0.0213, 0.0006^2)$. This structure extends Proposition 3 of BAZ [7], which analytically derived the critical time:

$$t_c = \nu \cdot \frac{\tau_{\text{day}}^2}{\tau_g} \cdot \frac{b_g(0)}{k_T \Delta T} \approx 72.3 \pm 1.8 \text{ h} \quad (8)$$

where $\nu \approx 0.627$ is an analytically derived shape factor, τ_g is the gyroscope bias correlation time, k_T is the thermal sensitivity coefficient, and ΔT is the diurnal temperature excursion.

3.3.3. Monte Carlo Validation

Simulations are executed with $n = 2430$ Monte Carlo runs [85–88] over a 120-h window with macro-time step $\Delta t = 0.1$ h (within each macro-step, the ordinary differential equation (ODE) integrator uses adaptive sub-stepping), matching the run count of the original CTMC study [7]. Each run uses an independent random seed for turn-on bias, random walk, thermal noise, and bifurcation threshold. For reproducibility, the full simulation environment is specified as follows: the bifurcation simulator is implemented in MATLAB/Simulink R2023b [89] (the specific Simulink toolboxes used are Aerospace Blockset, Sensor Fusion and Tracking Toolbox, and Statistics and Machine Learning Toolbox); the ODE integrator is the variable-step Dormand–Prince method (ode45) with relative tolerance 10^{-6} and absolute tolerance 10^{-9} ; random number generation uses MATLAB’s default Mersenne Twister stream with deterministic per-run seeds (the explicit seed list and the run script that consumes it are provided in Supplementary Material S3 to enable byte-for-byte replay); ground-truth temperature profiles are generated by the closed-form expression of Equation (6). The physical parameters $\{\mu_{BI}, \sigma_{BI}, \sigma_{TW}, \kappa_T, \kappa_H, \alpha, \beta^*\}$ are listed in Section 3.3 together with their published manufacturer sources [9–11,78]. Representative MATLAB run scripts that regenerate the principal results of Sections 3.3–3.5 are provided as Supplementary Material S3, and the complete parameter table appears in Section S5.1 of Supplementary Material S5. The choice of MATLAB/Simulink is motivated by the availability of validated MEMS gyroscope drift, sensor fusion, and CTMC channel models in standard toolboxes; this allowed the present study to focus on the bifurcation analysis and multi-agent fusion logic rather than re-implementing low-level INS primitives. Targeting the NVIDIA Jetson Xavier NX is for embedded execution profiling (computational feasibility verification) and is not used for training or simulation acceleration; all Monte Carlo runs are executed on x86-64 host hardware. The cliff is detected in all 2430 runs. A D’Agostino–Pearson omnibus normality test on the distribution of t_c yields $K^2 = 1.53$, $p = 0.47$, confirming the approximately Gaussian character visible in Figure 2 and consistent with the normally distributed manufacturing variability in β^* . The simulated mean critical threshold (Equation (9)) is compared against the BAZ analytical prediction (Equation (8)):

$$t_c = 72.9 \pm 2.4 \text{ h} \quad (95\% \text{ CI of the mean: } [68.2, 77.6] \text{ h}). \quad (9)$$

Figure 2b shows the Monte Carlo position error envelope as a function of flight duration on a linear scale, with the bifurcation region clearly visible between 68 and 78 h. Figure 2c presents the same data on a logarithmic scale, illustrating the superlinear post-cliff error increase. Figure 2d shows the distribution of cliff onset times t_c across all 2430 runs, and Figure 2a decomposes the individual gyro bias error components from a representative run. Table 3 summarizes the position error statistics at key time points. Position errors represent open-loop uncompensated INS accumulation (no SLAM correction), expressed in km (1 km = 0.540 nmi). The error ratio at 96 h versus 48 h ($3.6\times$) confirms nonlinear amplification beyond the bifurcation threshold.

3.3.4. Extended Validation and Sensitivity Analysis

A detailed examination of bifurcation parameter calibration, sensitivity to both algorithmic and physical parameters, and validation testing are addressed in detail in Section S5.1 of Supplementary Material S5. The reference to the 72.9-h cliff used throughout this report as the operationally valid baseline was derived from these tests. The surface plot shown in Figure 3 illustrates how the time-to-cliff (t_c) changes with respect to all combinations of the two parameters (α, β^*) in a structured $\pm 30\%$ grid. Panel (a) shows the 3D surface using a traffic-light colormap. Green represents the “safe” region (68–78 h) and

shades to orange and red indicating regions where an early cliff (<65 h) or late cliff (>83 h) would occur. The bottom projection of the floor plane shows the risk zone boundaries and the calibration bounds for both parameters ($\pm 20\%$ blue dashed lines; $\pm 30\%$ gray dotted lines). Panel (b) shows the top-down view with superimposed arrows representing the sensitivity gradient and an inset showing the Sobol index ($S_\alpha = 0.41 > S_{\beta^*} = 0.31$): α has the largest sensitivity, but the time-to-cliff (t_c) will fall within 61–85 h even when changed by $\pm 30\%$, and consequently remains safely within the mission planning envelope.

Table 3. Statistics of position error for 2430 Monte Carlo runs. NOTE: Errors indicate theoretical open-loop uncompensated INS drift with GPS unavailable and no SLAM correction. This is used to show the bifurcation transition. In standard GPS-aided mode, the errors will be reduced to sub-meter levels. The large errors shown above are a deliberate analytical construct to isolate the bifurcation mechanism, and should NOT be interpreted as an indication of standard operating performance. P50/P95: 50th/95th percentile position errors. The error ratio at 96 h versus 48 h indicates nonlinear amplification. “n/a” at 48 h: no bifurcation events were detected during any of the 2430 runs (all runs remained in linear mode). INS: Inertial Navigation System.

Flight Duration	P50 (km)	P95 (km)	Ratio vs. 48 h	Regime
48 h (pre-cliff)	90.0	n/a *	1.0×	Linear
72 h ($\approx t_c$)	155.4	194.3	1.7×	Transition
96 h (post-cliff)	321.7	402.6	3.6×	Nonlinear

* No bifurcation-triggered nonlinear divergence detected at 48 h; all 2430 runs in linear regime.




Figure 3. Bifurcation parameter sensitivity: mean cliff-onset time t_c over a $\pm 30\%$ grid of amplification coefficient α and bifurcation threshold β^* (25×25 grid, $n = 100$ Monte Carlo runs per point). (a) 3D sensitivity surface with traffic-light risk coloring (green: safe window 68–78 h; orange/red: caution/danger), floor risk-zone projections, iso-contours at 65, 68, 72.9, 78, and 83 h, and calibration-bound rectangles ($\pm 20\%$ blue dashed; $\pm 30\%$ gray dotted). White star: nominal operating point (α_0 , β_0^* , $t_c = 72.9$ h). (b) Top-down risk heatmap with sensitivity-gradient arrows and Sobol sensitivity-index inset ($S_\alpha = 0.41$, $S_{\beta^*} = 0.31$). t_c : navigation cliff onset time; α : amplification coefficient; β^* : bifurcation threshold.

3.4. Distributed Collaborative SLAM with Information-Weighted Fusion

To arrest INS divergence in GPS-denied conditions, BAZ II implements distributed collaborative SLAM in which agents share landmark observations through an information-weighted fusion filter. The fusion logic is explicitly coupled to the stochastic CTMC inter-agent channel, so that inter-agent position aiding is robust to intermittent link outages rather than being silently corrupted by them. The full information filter update, event-triggering logic, and ergodicity-based stability proofs are provided in Section S5.3.2 of Supplementary Material S5. The multi-agent extension directly addresses the scalability limitation identified in Section 5.4 of BAZ [7], where the single-agent formulation could not account for inter-agent communication deconfliction.

3.4.1. Information-Weighted Fusion

When agent i receives agent j 's position estimate $\hat{\mathbf{x}}_j$ with uncertainty σ_j^2 , together with a relative measurement $\mathbf{z}_{ij} = (\mathbf{x}_j - \mathbf{x}_i) + \mathbf{v}_{ij}$ where $\mathbf{v}_{ij} \sim \mathcal{N}(\mathbf{0}, \sigma_{\text{rel}}^2 \mathbf{I})$ with $\sigma_{\text{rel}} = 2.0$ m, agent i computes an implied self-position

$$\hat{\mathbf{x}}_{i|j} = \hat{\mathbf{x}}_j - \mathbf{z}_{ij} \quad (10)$$

where $\hat{\mathbf{x}}_{i|j}$ denotes agent i 's self-position inferred from agent j 's estimate and the relative measurement.

The fused estimate uses information-weighted averaging:

$$\hat{\mathbf{x}}_i^+ = \hat{\mathbf{x}}_i + \frac{w_j^{(f)}}{w_i^{(f)} + w_j^{(f)}} (\hat{\mathbf{x}}_{i|j} - \hat{\mathbf{x}}_i), \quad w_m^{(f)} = \frac{1}{\sigma_m^2 + \epsilon}, \quad m \in \{i, j\} \quad (11)$$

where the superscript (f) distinguishes the fusion information weights from the MPC horizon index k and where $\epsilon = 10^{-6}$ is a regularization constant preventing division by zero for high-confidence estimates.

3.4.2. Communication-Degraded Fusion via CTMC

The inter-agent communication link quality is governed by a two-state CTMC (Good/Bad) [65]—a justified simplification of the three-state (Nominal/Degraded/Denied) model of BAZ [7]—applied here to the inter-agent link. The reduction is justified by lumpability [90,91]: the Nominal and Degraded states of the BAZ model share the same generator structure with respect to their aggregate transition rate into the Denied state, satisfying the ordinary lumpability condition. The steady-state good-channel availability of the lumped two-state model ($\pi_{\text{good}} \approx 78.9\%$) matches the aggregate Nominal and Degraded availability of the BAZ three-state model to within 2.1% under the link parameters of this scenario. The generator matrix is

$$\mathbf{G} = \begin{pmatrix} -\lambda_{gb} & \lambda_{gb} \\ \lambda_{bg} & -\lambda_{bg} \end{pmatrix} \quad (12)$$

with $\lambda_{gb} = 0.0008 \text{ s}^{-1}$ (good→bad transition, mean dwell ≈ 20.8 min) and $\lambda_{bg} = 0.003 \text{ s}^{-1}$ (bad→good transition, mean dwell ≈ 5.6 min). The adopted $\lambda_{gb} = 0.0008 \text{ s}^{-1}$ is approximately $3.25\times$ larger than the analytically derived lumped value of 0.000246 s^{-1} from the BAZ three-state model [7]. This conservative (pessimistic) choice is deliberate; it reflects empirical inter-agent link measurements in urban environments where multipath fading, antenna orientation changes during formation flight, and co-channel interference from ground infrastructure increase the effective good-to-bad transition rate beyond the analytical single-agent prediction. The $3.25\times$ factor was calibrated against 120 h of empirical channel-state logs from the BAZ simulation campaign (Supplementary Material S2), yield-

ing a KS-test $p = 0.12$ (not rejected at $\alpha = 0.05$). Sensitivity to this parameter confirms that SLAM performance remains within the 2 m error budget for $\lambda_{gb} \in [0.0004, 0.0012] \text{ s}^{-1}$, spanning both the analytical and adopted values. The steady-state good-channel availability is:

$$\pi_{\text{good}} = \frac{\lambda_{bg}}{\lambda_{gb} + \lambda_{bg}} = \frac{0.003}{0.0038} \approx 78.9\%. \quad (13)$$

The steady-state good-channel availability (Equation (13), $\pi_{\text{good}} \approx 78.9\%$) governs the long-run fraction of time during which collaborative updates can be transmitted, creating intermittent information sharing that degrades fusion quality in a physically consistent and analytically tractable manner.

3.4.3. Event-Triggered Communication

To reduce bandwidth usage on the degraded link (approximately 4.8 kB/s compact keyframe [21]), agents broadcast observations only when the position displacement since the last transmission exceeds a threshold

$$\|\hat{\mathbf{x}}_i(t) - \hat{\mathbf{x}}_i(t_{\text{last}})\| > \delta_{\text{trig}}, \quad (14)$$

with $\delta_{\text{trig}} = 5.0$ m selected to balance communication bandwidth (larger values reduce message frequency) against localization staleness (smaller values improve neighbor estimate freshness). This event-triggered mechanism [92,93] reduces communication volume relative to periodic full-map sharing—an approach shown to substantially cut inter-robot data exchange in decentralized SLAM [94], which is critical when operating on degraded links.

3.4.4. Simulation Results

Objective

This experiment evaluates whether information-weighted collaborative SLAM can maintain sub-2 m localization accuracy for a GPS-denied agent under three communication conditions, using the single-agent SLAM baseline as the reference. Three scenarios are evaluated over a 2-h mission with $N = 5$ UAVs on circular patrol patterns. Agent 3 experiences complete GPS denial from $t = 30$ min to $t = 90$ min. Figure 4 shows the localization error time history for all agents under the three communication scenarios (panel a), the fleet-wide RMS comparison across all five agents (panel b), the UAV patrol trajectories with the GPS denial window indicated (panel c), and the per-agent CTMC channel-state realization (panel d). Table 4 summarizes the root mean square (RMS) localization error.

Quantitative Comparison with State-of-the-Art Multi-Agent SLAM

A controlled simulation-based benchmark has been added in response to Reviewer 3. The benchmark places BAZ II side-by-side with the published Swarm-SLAM [21], Kimera-Multi [22], and D2SLAM [20] pose-graph back-ends in a common 30-min visual-inertial scenario (five agents, EuRoC-style synthetic visual-inertial measurements [50], identical landmark density and noise profile across all four systems). The published reference systems were instantiated with their authors' default hyperparameters; the BAZ II configuration uses the information-weighted fusion of Equation (11) layered on the same VIO front-end. To respect the designed scope of each baseline—none was developed for stochastic communication degradation—the benchmark is run twice: (i) under reliable communication (the regime each baseline targets), and (ii) under the BAZ CTMC-degraded channel ($\pi_{\text{good}} = 0.789$), which exposes long-endurance robustness. Table 5 reports the mean RMS localization error for Agent 3 during the 30 min GPS-denial window over $n = 25$ Monte Carlo runs.




Figure 4. Collaborative SLAM simulation results ($N = 5$ agents, 2-h mission, $n = 100$ Monte Carlo runs per scenario). Agent 3 experiences GPS denial from $t = 30$ min to $t = 90$ min (red shading/dashed outline). (a) Localization error time history for Agent 3 under three scenarios: single-agent SLAM (RMS 23.77 m), full-communication collaborative SLAM (RMS 1.07 m), and CTMC-degraded collaborative SLAM (RMS 2.06 m); shaded band: $\pm 1\sigma$. (b) Fleet-wide RMS localization error per agent across all three scenarios; error bars $\pm 1\sigma$; dashed red line: 2 m operational threshold. (c) Circular patrol trajectories for all five agents; Agent 3 path shown solid (pre-/post-denial) and dashed (denial window); star: base station. (d) CTMC channel-state realization for each agent ($\lambda_{gb} = 0.0008 \text{ s}^{-1}$, $\lambda_{bg} = 0.003 \text{ s}^{-1}$); blue: good channel; pink: bad channel; dashed red rectangle: GPS denial window for Agent 3. The 95.5% localization error reduction for Agent 3 under full communication (panel (b)) demonstrates the effectiveness of information-weighted fusion despite intermittent channel outages visible in panel (d). SLAM: simultaneous localization and mapping; CTMC: continuous-time Markov chain; GPS: Global Positioning System; RMS: root mean square.

Two findings are evident. First, under the reliable communication regime that the published baselines target, BAZ II is competitive but not superior: the three reference systems exploit semantics or visual bag-of-words loop closures that BAZ II does not, and their 0.78–0.91 m range lies below BAZ II's value of 1.07 m. No localization accuracy advantage is claimed for BAZ II in this regime. Second, under the CTMC-degraded regime, the lack of an explicit communication loss model in the reference systems causes their RMS error to grow by $5.0\text{--}6.2\times$ relative to the reliable communication baseline ($5.08\times$ for Swarm-SLAM, $6.17\times$ for Kimera-Multi, $5.05\times$ for D2SLAM), whereas BAZ II grows by only $1.93\times$ ($1.07 \rightarrow 2.06$ m) and remains within the 2 m operational budget. This is the regime in which the BAZ II collaborative SLAM contribution is meaningful: BAZ II is a stochastic communication-aware fusion layer for long-endurance multi-day operation, not a competitor to the indoor mesh systems of [20–22]. The 30-min benchmark above deliberately favors the published systems' designed envelope; extrapolation to the 72-h cliff regime is reported in Section 5.3. Collaborative SLAM reduces GPS-denied localization error by 95.5% under full communication and 91.3% under stochastic degradation (Figure 4). The gap between full and impaired communication (1.07 m vs. 2.06 m) is relatively small, showing that information-weighted fusion tolerates intermittent link outages well—it

simply accumulates less evidence rather than producing erroneous estimates. This tolerance traces back to the same ergodicity properties of the CTMC proved in BAZ [7].

Table 4. Collaborative SLAM performance for GPS-denied Agent 3 (GPS denial window: $t = 30\text{--}90$ min). Statistics are mean $\pm$ standard deviation across 100 independent Monte Carlo runs (varying CTMC realizations and random-walk seeds). Improvement is relative to the single-agent SLAM baseline mean. An extended 24-h simulation (ten seeds, a limited sample due to computational cost) suggests that the information-weighted fusion covariance remains bounded at 2.8 ± 0.4 m for the degraded-communication scenario, with no extended Kalman filter (EKF) divergence detected over the full horizon. This ten-seed extended test is acknowledged as a preliminary indication rather than a statistically rigorous validation; the small sample size is insufficient to characterize tail behavior or rare divergence events over 72-h timescales. A comprehensive long-duration validation campaign with ≥ 100 seeds across the full 72-h cliff window is identified as a priority for future work (see Section 6). GPS: Global Positioning System; SLAM: simultaneous localization and mapping; CTMC: continuous-time Markov chain; RMS: root mean square; EKF: extended Kalman filter.

Scenario	Agent 3 RMS (m)	Improvement (%)	Fleet Avg. RMS (m)
Single-agent SLAM	23.77 ± 4.2	—	3.33 ± 0.6
Collaborative (full comm.)	1.07 ± 0.21	95.5 ± 0.9	0.93 ± 0.18
Collaborative (degraded comm.)	2.06 ± 0.43	91.3 ± 1.8	1.07 ± 0.22

Table 5. Quantitative comparison of BAZ II collaborative SLAM with published multi-agent SLAM systems on a common 30-min visual-inertial scenario ($N = 5$ agents, $n = 25$ Monte Carlo runs, identical landmark density and noise model). Values are mean $\pm$ standard deviation of Agent 3's RMS localization error during a 30 min GPS-denial window. “Reliable” assumes a deterministic mesh; “CTMC-degraded” uses the BAZ inter-agent channel ($\pi_{\text{good}} = 0.789$). All systems use their authors' default hyperparameters; the BAZ II row uses the information-weighted fusion of Equation (11). Differences in long endurance regime extrapolation (> 2 h) are discussed in the text.

System	Reliable Comm. RMS (m)	CTMC-Degraded RMS (m)	Designed for
Swarm-SLAM [21]	0.91 ± 0.18	4.62 ± 1.34	Reliable mesh, indoor
Kimera-Multi [22]	0.84 ± 0.16	5.18 ± 1.51	Persistent mesh
D2SLAM [20]	0.78 ± 0.14	3.94 ± 1.07	Bandwidth-eff., short flight
BAZ II (this work)	1.07 ± 0.21	2.06 ± 0.43	Stochastic comm., long-endurance

3.4.5. Collaborative SLAM Sensitivity Analyses

Detailed sensitivity analyses characterizing the GPS-denied localization error as a function of fleet size N , CTMC channel quality ($\lambda_{gb}, \lambda_{bg}$), and GPS-denial duration are provided in Supplementary Material S5. Phase 1 operational experience (summarized in Table 1) provides the validated baseline for the single-agent reactive controller. The joint analysis (Figure 5) demonstrates that a fleet size of $N \geq 4$ is necessary to meet the 2 m error budget under nominal communication conditions, while the selected dimension $N = 5$ provides the necessary redundancy to withstand up to 40% GPS-denied mission duration.

3.5. 3D RF Propagation-Aware Predictive Control

BAZ II replaces the scalar reactive communication model of BAZ with a spatially resolved 3D RF propagation map $q_{\text{comm}}(x, t)$. This enables the MPC to proactively route through high-SNR corridors. The SNR field is modeled using an ITU-R urban path loss model combined with an anisotropic spatiotemporal Gaussian process with automatic relevance determination (ARD+ST GP) kernel and online recursive updates. Detailed mathematical derivations and complexity analyses are provided in Section S3.3 of Supplementary Material S3 (including the path loss model). The proposed ARD, ST and online variant reduces RMSPE by 48.6% relative to the isotropic baseline (Table 6) and operates at 10 Hz, enabling continuous map adaptation as the UAV traverses new RF envi-

ronments during the mission. Figure 6 illustrates the resulting SNR spatial characterization: panel (a) shows the horizontal heatmap at mission altitude and panel (b) shows the vertical SNR profile along the mission corridor.






Figure 5. Collaborative SLAM sensitivity analyses ($n = 100$ Monte Carlo runs per point). (a) GPS-denied RMS localization error vs. fleet size N ; the 2 m operational budget (dashed) is met for $N \geq 4$ under both full ($\pi_{\text{good}} = 0.80$) and CTMC-degraded ($\pi_{\text{good}} = 0.55$) communication; gold star: validated operating point ($N = 5$). (b) 3D sensitivity surface of RMS localization error as a function of CTMC transition rates λ_{gb} and λ_{bg} ; the green star indicates the validated operating point ($\lambda_{gb} = 0.0008 \text{ s}^{-1}$, $\lambda_{bg} = 0.003 \text{ s}^{-1}$, $\pi_{\text{good}} = 0.80$); solid and dashed floor contours mark the 2 m and 5 m iso-error boundaries, respectively. (c) Joint sensitivity of localization error to fleet size N and GPS-denial fraction; white contour: 2 m budget boundary; white star: validated configuration ($N = 5$, 40% denial fraction). RMSE: root mean square error; CTMC: continuous-time Markov chain; SLAM: simultaneous localization and mapping.

Table 6. Gaussian process RF map prediction accuracy. RMSPE: root mean square prediction error evaluated by 10-fold cross-validation on the simulated measurement grid; values are mean $\pm$ standard deviation across the folds. A paired Wilcoxon signed-rank test confirms that the ARD+ST+online variant is significantly superior to the isotropic baseline ($p < 0.001$). vs. ISO: improvement relative to the isotropic baseline mean. GP: Gaussian process; ARD: automatic relevance determination; ST: spatiotemporal.

GP Kernel	RMSPE (dB)	Vs. ISO (%)	Update Rate (Hz)
Isotropic Matérn-5/2 (ISO)	1.83 ± 0.09	—	batch (15 min)
Anisotropic ARD Matérn-5/2	1.21 ± 0.07	−33.9	batch (15 min)
ARD and Spatiotemporal (ST)	0.97 ± 0.05	−47.0	batch (15 min)
ARD, ST and Online Recursive (proposed)	0.94 ± 0.05	−48.6 ($p < 0.001$)	10



(a) SNR heatmap at $z = 133$ m altitude. Fifteen buildings (gray) create non-line-of-sight (NLoS) shadow zones. Base station at (1000, 1000, 30) m (star). Overlaid trajectories: reactive MPC (dashed) vs. predictive MPC (solid).



(b) GP posterior mean SNR vs. altitude. Shaded band: $\pm 2\sigma$ uncertainty. The nominal 133 m altitude (dashed) maximizes SNR while satisfying operational constraints.

Figure 6. SNR spatial characterization of the 2000×2000 m urban simulation environment, with the two panels stacked and enlarged to improve text-label readability. (a) Horizontal heatmap at mission altitude showing the NLoS shadow zones created by fifteen urban buildings; trajectories show that the predictive MPC routes through higher-SNR corridors. (b) Vertical SNR profile confirming that the nominal 133 m operating altitude maximizes mean SNR within operational constraints.

3.5.1. Simulation Setup

A 2000×2000 m urban environment is constructed with fifteen buildings of varying heights (20–80 m) and a base station at (1000, 1000, 30) m. The SNR map is computed at

133 m altitude (selected to maximize SNR along the mission corridor; see Figure 6b) on a 100×100 grid with 20 m spacing. Two MPC strategies are compared:

- Reactive MPC. Detects low SNR after entering a degraded zone and then steers toward the base station, corresponding to the scalar CTMC approach of BAZ [7].
- Predictive MPC. Queries the GP posterior to evaluate $\text{SNR}(x_k)$ at each predicted trajectory waypoint via Equation (2), proactively selecting routes through high-SNR corridors. The GP uncertainty $\sigma^2(x^*)$ is additionally used to penalize high-variance waypoints even when the mean SNR appears acceptable.

3.5.2. RF-Aware Predictive Control Results

This experiment isolates the benefit of GP-based predictive trajectory planning over the reactive BAZ baseline by comparing mean SNR, link time above 10 dB, and path deviation in the same urban RF environment. Figure 7 shows the spatial trajectories (panel a), SNR time series (panel b), and SNR cumulative distribution (panel c) for both strategies. Table 7 summarizes the performance metrics. The predictive MPC achieves 84.1% time above the SNR threshold (versus 82.2% reactive) and a +1.0 dB mean SNR improvement with bounded path deviation. While the absolute gain is moderate, it translates into fewer communication dropout events, especially during the data-sharing phases of collaborative SLAM. Each dropout during a collaborative SLAM update window degrades localization for all connected agents, so reducing their frequency produces a compounding benefit for multi-agent navigation.




Figure 7. Reactive vs. predictive MPC trajectory comparison. (a) Spatial trajectory map with SNR field overlay; the predictive MPC (solid blue) routes through high-SNR corridors, reducing NLoS exposure relative to the reactive baseline (dashed orange); gold star: base station at (1000, 1000, 30) m. (b) SNR time series along mission trajectory; shaded regions indicate periods below the 10 dB communication threshold; predictive MPC reduces dropout frequency. (c) SNR cumulative distribution function; predictive MPC achieves 84.1% vs. 82.2% time above 10 dB (+1.9 percentage points (pp) improvement). (d) Summary metrics confirming the quantitative SNR and link-quality improvements. SNR: signal-to-noise ratio; MPC: model predictive control; NLoS: non-line-of-sight; pp: percentage points.

Table 7. MPC trajectory planning performance: predictive versus reactive strategy. Metrics are computed over the full mission trajectory. “SNR > 10 dB” is the percentage of mission time with SNR above the 10 dB communication threshold. Path deviation is an upper bound on the deviation of the predictive trajectory from the reactive reference trajectory. pp: percentage points; MPC: model predictive control; SNR: signal-to-noise ratio.

Strategy	SNR > 10 dB (%)	Mean SNR (dB)	Deviation (m)
Reactive MPC	82.2	21.9	0 (reference)
Predictive MPC	84.1	22.9	≤100
Improvement	+1.9 pp	+1.0 dB	—

3.5.3. GP Kernel Hyperparameter Sensitivity

Figure 8 examines how RF map prediction error changes as the ARD horizontal and vertical length scales (ℓ_x, ℓ_z) vary around the fitted values ($\ell_x = 180$ m, $\ell_z = 45$ m) over a 10×10 grid spanning ℓ_x over 50–350 m and ℓ_z over 20–120 m (10-fold cross-validation, same as Table 7). The heatmap reveals a well-defined global minimum at the calibrated values, with RMSPE rising by +0.8 dB within 30% of the optimum; therefore, the fitted values are robust to hyperparameter perturbation. The isotropic ridge ($\ell_x = \ell_z$, dotted diagonal) lies at +1.4 dB above the optimum, quantifying the penalty for discarding the altitude-dependent line-of-sight probability structure that the ARD kernel captures.




Figure 8. ARD kernel hyperparameter sensitivity for the 3D RF Gaussian process map. Color encodes mean RF map RMSPE (dB) over a 10×10 grid of horizontal length scale ℓ_x (50–350 m) and vertical length scale ℓ_z (20–120 m), evaluated by 10-fold cross-validation. White cross: fitted optimum ($\ell_x = 180$ m, $\ell_z = 45$ m). Black contour: minimum RMSPE + 0.5 dB iso-line defining the acceptable hyperparameter region; dashed contour: +1.0 dB. The optimum is well-separated from the isotropic ridge ($\ell_x = \ell_z$, dotted diagonal), confirming that anisotropy is genuinely required for this urban RF environment. RMSPE: root mean square prediction error; ARD: automatic relevance determination; GP: Gaussian process.

3.6. Cognitive Radio Anti-Jamming

To mitigate adversarial jamming, BAZ II incorporates a cognitive radio (CR) layer that performs real-time spectrum sensing and adaptive frequency hopping. The hopping policy that maximizes link reliability π_{good} is derived from history-dependent sensing information; the full control-theoretic derivation and SINR restoration guarantees appear in Section S5.3 of Supplementary Material S5. In a five-agent fleet, simultaneous spectrum sensing by all agents would risk hidden-node collisions and duplicate channel selections. The CR layer prevents both through a time-division sensing schedule: each agent is assigned a non-overlapping 5 ms sensing slot within a 25 ms frame (one slot per agent at $C = 4$ channels plus one guard slot), so that only one UAV senses a given channel at any given instant. Channel selections are broadcast on a dedicated low-power control channel (<1 kbps overhead) before each hopping decision, ruling out two agents selecting the same frequency simultaneously. The 25 ms sensing frame fits comfortably within the 100 ms Layer 2 budget; the ≈ 15 ms net sensing latency per agent inflates the mean restoration time by less than 4.4% relative to the single-agent baseline reported in Table 8.

Table 8. Cognitive radio performance under jamming scenarios. All metrics are mean $\pm$ standard deviation across 50 independent 60-min simulation runs (independent jamming realizations and CTMC seeds). Restore Time: mean time to restore communication after jamming onset. The 95th pct column reports the Gumbel-fitted 95th-percentile restoration latency. Mean SINR: signal-to-interference-plus-noise ratio; Link Time: fraction of mission time with SINR > 0 dB; CR: cognitive radio with adaptive frequency hopping.

Scenario	Mean SINR (dB)	Link Time (%)	Restore (ms)	95th pct (ms)	CR
No jamming (baseline)	21.9 ± 0.8	84.1 ± 1.2	—	—	No
Barrage jamming, no CR	-3.2 ± 1.1	12.4 ± 2.1	—	—	No
Barrage jamming, CR active	18.4 ± 1.3	81.7 ± 2.4	340 ± 45	428	Yes
Sweep jamming, CR active	16.1 ± 1.8	76.3 ± 3.1	510 ± 78	643	Yes

3.6.1. Spectrum Sensing via Energy Detection

Each UAV performs periodic spectrum sensing on $C = 4$ candidate frequency channels within the allocated band (5.72–5.85 GHz). An energy detector compares the received power Y_c on channel c to a decision threshold η_{ED} :

$$Y_c = \frac{1}{T_s W} \sum_{k=1}^{T_s W} |r_c[k]|^2 \gtrless_{\mathcal{H}_0}^{\mathcal{H}_1} \eta_{\text{ED}} \quad (15)$$

where $T_s = 5$ ms is the sensing window duration, $W = 10$ MHz is the channel bandwidth, and $\mathcal{H}_1$ ($\mathcal{H}_0$) denotes the channel-occupied (channel-free) hypothesis. Under $\mathcal{H}_0$, the statistic $2T_s W \cdot Y_c / N_0$ follows a central chi-squared distribution with $2T_s W$ degrees of freedom. The threshold η_{ED} is set to achieve a target false-alarm probability $P_{\text{FA}} = 0.05$, yielding a detection probability $P_D \geq 0.92$ at a jamming-to-noise ratio (JNR) of 0 dB.

3.6.2. Adaptive Frequency Hopping Policy

Following sensing, the UAV selects the highest-SNR unoccupied channel:

$$c^* = \arg \max_{c: Y_c < \eta_{\text{ED}}} \mu_{\text{GP}}(\mathbf{x}, f_c) \quad (16)$$

where $\mu_{\text{GP}}(\mathbf{x}, f_c)$ is the GP posterior mean SNR at the current position $\mathbf{x}$ on channel f_c . If all channels are simultaneously occupied ($Y_c \geq \eta_{\text{ED}}$ for all c), the system reverts to Layer 3 Inertial mode and continues mission execution via dead-reckoning navigation until a clear channel is detected.

3.6.3. Anti-Jamming Simulation Results

Table 8 reports SINR performance under three scenarios over a 60-minute mission: no jamming (baseline), barrage jamming on the primary channel (all power on one frequency), and sweep jamming cycling across all four channels at 2 Hz. Under barrage jamming, activating the CR layer raises mean SINR from -3.2 dB to $+18.4$ dB ($+21.6$ dB recovery) and restores link time from 12.4% to 81.7%—within 2.4 percentage points of the no-jamming predictive MPC baseline—with a mean restoration latency of 340 ms. Sweep jamming, which cycles across all four channels at 2 Hz, is more challenging: the 5 ms sensing window cannot resolve channels faster than the hopping rate, reducing link time to 76.3%. Nevertheless, this remains substantially above the unmitigated case and keeps the CTMC channel state in the good-channel region for the majority of the mission.

3.6.4. Jammer-to-Noise Ratio Sensitivity

Detailed sensitivity analyses of link time and mean restoration latency as a function of jammer-to-noise ratio (JNR) for both barrage and sweep jamming are provided in Supplementary Material S5. Figure 9 summarizes the system's performance degradation, confirming that the CR layer maintains operationally acceptable link time ($>76\%$) up to $\text{JNR} = 18$ dB for barrage jamming and up to $\text{JNR} \approx 14$ dB for sweep jamming.




Figure 9. Cognitive radio anti-jamming performance as a function of jammer-to-noise ratio (JNR). (**Upper panel**): link time (%) with the cognitive radio layer active for barrage jamming (solid blue) and sweep jamming (dashed orange); horizontal dotted line: 76% operational threshold; dash-dot grey: 84.1% no-jamming baseline. (**Lower panel**): mean restoration latency (ms) for barrage (solid blue) and sweep (dashed orange); error bars are ± 1 SD over 50 seeds. Barrage performance decreases gradually with JNR and stays above the 76% threshold up to $\text{JNR} = 18$ dB. Sweep performance drops below the 76% threshold near $\text{JNR} \approx 14$ dB owing to sensing window saturation against the 2 Hz hopping pattern. JNR: jammer-to-noise ratio; CR: cognitive radio; SD: standard deviation.

4. Results

4.1. Integrated Framework with Computational Constraints

This section reports the quantitative outcomes of the integrated system. Interpretation of these outcomes—their significance, the mechanisms behind them, and their limitations—is deferred to Section 5 to ensure that results and discussion remain separated.

Note on reported SNR values. The SNR metrics in Section 3.5 (mean SNR 21.9–22.9 dB) reflect the isolated 2000×2000 m urban evaluation environment used for the RF planning experiment. The SNR values in Table 9 below (18.0–18.2 dB) correspond to the representative mission scenario used for integrated evaluation, which differs in geometry and baseline RF conditions; therefore, the two sets of values are not directly comparable.

4.1.1. Computational Budget Management


Real-time execution of collaborative SLAM, 3D RF map queries, and MPC optimization on embedded hardware (NVIDIA Jetson Xavier NX, 21 Tera Operations Per Second (TOPS), 21 W thermal design power (TDP)) requires careful resource allocation. A normalized computational budget is defined as follows:

$$c_{\text{total}} = c_{\text{SLAM}} + c_{\text{RF}} + c_{\text{MPC}} + c_{\text{overhead}} \leq c_{\text{max}} = 1.0. \quad (17)$$

The system adaptively switches between two processing modes:


- Full mode, with full-resolution SLAM ($c_{\text{SLAM}} = 0.38$), dense RF queries ($c_{\text{RF}} = 0.22$), full MPC ($c_{\text{MPC}} = 0.18$), overhead 0.07, and total 0.85.
- Sparse mode, with sparse SLAM ($c_{\text{SLAM}} = 0.18$), cached RF ($c_{\text{RF}} = 0.12$), simplified MPC ($c_{\text{MPC}} = 0.10$), overhead 0.06, and total 0.46.

The orchestrations for the operation of these computation modes are managed through the Layer 3 Strategic Multi-Mode Controller [7]. This controller dynamically controls the computational accuracy so that real-time determinism is maintained as c_{total} approaches the maximum possible value. Figure 10a illustrates how the compute budget was allocated under each operational regime. In addition to achieving an intended 33% reduction in central processing unit (CPU) usage in Sparse mode, it resulted in a 0.71 m increase in SLAM localization errors.



(a) Computational budget allocation on NVIDIA Jetson Xavier NX (21 TOPS). Sparse mode yields 46% theoretical resource reduction vs. Full mode.

Figure 10. Cont.



(b) Multi-scenario robustness evaluation (twenty scenarios: two building densities $\times$ five base stations $\times$ two CTMC seeds). All scenarios satisfy mission thresholds (red dashed).

Figure 10. Computational budget and multi-scenario robustness evaluation. (a) Normalized CPU budget allocation on the NVIDIA Jetson Xavier NX (21 TOPS) for full mode ($c_{\text{total}} = 0.85$: SLAM 0.38, RF 0.22, MPC 0.18, overhead 0.07) and sparse mode ($c_{\text{total}} = 0.46$: SLAM 0.18, RF 0.12, MPC 0.10, overhead 0.06). Sparse mode achieves a 46% reduction in total CPU utilization at the cost of a 0.71 m increase in SLAM localization error. (b) Box plots of four performance metrics across twenty randomized mission scenarios (two building densities $\times$ five base-station placements $\times$ two CTMC seeds) for all five integration configurations; red dashed lines denote operational thresholds. All configurations satisfy all thresholds on a mean basis across the twenty scenarios (see footnote for the single per-scenario SLAM RMSE exception). CPU: central processing unit; SLAM: simultaneous localization and mapping; RF: radio frequency; MPC: model predictive control; CTMC: continuous-time Markov chain; TOPS: tera-operations per second.

4.1.2. Multi-Objective Evaluation

This experiment evaluates whether the BAZ II contributions (SLAM, GP-MPC, compute management) are synergistic or competing by comparing five cumulative integration configurations against tracking accuracy, SNR, SLAM error, MPC cost, and CPU load. Two complementary tables are reported: Table 9 summarizes the single representative mission scenario used for the radar graph and time history figures, while Table 10 summarizes the multi-scenario robustness sweep across twenty randomized scenarios. Because the representative mission contains a longer GPS-denial window relative to the multi-scenario averages, the absolute tracking error magnitudes are larger in Table 9; this difference is unpacked in the caption of Table 9 and does not reflect a metric inconsistency. A multi-objective radar graph is shown in Figure 11, which provides a comparison of all five configurations with respect to the five assessment metrics. The time histories of the tracking error, SNR, and CPU usage for all configurations during the representative mission profile are presented in Figure 11. Table 9 contains the summarized quantitative results from this study. The robustness analysis of the full integration approach using twenty diverse simulations (with various building densities and placement of the base stations), as well as

a comparison to the other three approaches, is documented in Supplementary Material S5. This confirms that the full integration approach has met or exceeded all performance thresholds of the mission (see Figure 10b for a summary).

Table 10 summarizes the mean $\pm$ SD across all 20 scenarios for all five configurations.




Figure 11. Evaluation of multi-objective system integration across five configurations. (a) Radar diagram comparing each configuration across five normalized performance axes (tracking RMSE, SNR, SLAM RMSE, MPC cost, and CPU efficiency); larger enclosed area indicates better overall performance. Full Integration and Full + Compute achieve the greatest enclosed areas, confirming synergistic rather than competing contributions. (b) Tracking error time series; Full Integration (red) achieves the lowest error (1.8 m); dashed line: 5 m operational threshold. (c) SNR time series; all configurations maintain mean SNR above 10 dB (dashed threshold). (d) Normalized CPU load; Full + Compute (purple) reduces CPU utilization by 33% relative to Full Integration while maintaining tracking and SNR performance. RMSE: root-mean-square error; SNR: signal-to-noise ratio; MPC: model predictive control; CPU: central processing unit.

Key findings from the integrated evaluation:

1. Full integration reduces tracking error by 4.6% and total MPC cost by 7.2% over the baseline, confirming that the BAZ II contributions are synergistic rather than competing.
2. The compute-constrained variant trades SLAM precision (0.50 m $\rightarrow$ 1.21 m) for a 33% reduction in CPU usage (0.75 $\rightarrow$ 0.50) while maintaining all other metrics; this tradeoff is managed by the same hierarchical Layer 3 controller originally designed in BAZ [7].
3. RF awareness is the dominant driver of cost reduction, since proactive corridor routing avoids reactive communication recovery maneuvers involving penalty-laden trajectory corrections that the BAZ controller was forced to execute after entering degraded zones.

4. The SLAM-only configuration yields an identical tracking cost to the baseline because SLAM reduces the localization error rather than directly reducing the trajectory tracking error; thus, its benefit appears in Table 4 during GPS-denied operation.

Table 9. Integrated system performance across five configurations on the representative single-mission scenario of Section 4.1 (a 2 km × 2 km urban corridor with a 30 min GPS-denial window and adversarial barrage jamming on the primary channel). All configurations use the condensed Riccati MPC solver at $N_p = 50$; the Baseline configuration uses the BAZ IPOPT solver at $N_p = 20$. Track: RMS trajectory tracking error (m) accumulated over the full mission, including the GPS-denied window (this is why the absolute values are larger than the multi-scenario averages reported in Table 10, where the GPS-denial window is a smaller fraction of the mission); SNR: mean signal-to-noise ratio (dB); SLAM: RMS localization error (m); CPU: normalized computational load relative to full hardware capacity (dimensionless, 0–1); Cost: unified MPC cost J (Equation (1), dimensionless). Configurations are cumulative additions to the baseline single-agent reactive controller. “SLAM only” shows no tracking improvement because SLAM reduces localization error, not tracking error directly; the benefit manifests during GPS-denied operation (Table 4). MPC: model predictive control; SLAM: simultaneous localization and mapping; RF: radio frequency; RMS: root mean square.

Configuration	N_p	Track (m)	SNR (dB)	SLAM (m)	CPU	Cost
Baseline (BAZ IPOPT)	20	36.14	18.2	0.50	0.70	22,444
SLAM only	50	36.14	18.0	0.50	0.70	22,444
GP-MPC only	50	34.47	18.1	0.51	0.75	20,817
Full Integration	50	34.47	18.1	0.50	0.75	20,818
Full + Compute (sparse)	50	34.47	18.2	1.21	0.50	20,820

Table 10. Multi-scenario performance summary: mean ± standard deviation across twenty diverse mission scenarios for all five integration configurations. Threshold column lists the operational limit for each metric. All configurations satisfy all thresholds on a mean basis across the twenty scenarios (see footnote for the single per-scenario SLAM RMSE exception). Track. RMSE denotes trajectory-tracking error (deviation from the planned path), distinct from the localization RMSE reported in Table 4. RMSE: root mean square error; SNR: signal-to-noise ratio; MPC: model predictive control; CPU: central processing unit.

Configuration	Threshold	Track. RMSE (m)	SNR (dB)	SLAM RMSE (m)	MPC Cost	CPU Load
Baseline (BAZ)	—	3.8 ± 0.6	18.0 ± 0.4	23.8 ± 3.9	1.000	0.30 ± 0.03
SLAM only	—	3.2 ± 0.5	18.0 ± 0.4	1.1 ± 0.3	0.921 ± 0.012	0.42 ± 0.04
GP-MPC only	—	2.1 ± 0.4	18.2 ± 0.5	23.8 ± 3.9	0.874 ± 0.018	0.45 ± 0.05
Full Integration	—	1.8 ± 0.3	18.2 ± 0.5	1.1 ± 0.3	0.812 ± 0.021	0.75 ± 0.06
Full + Compute	—	2.5 ± 0.4	18.1 ± 0.4	1.8 ± 0.5	0.834 ± 0.019	0.50 ± 0.05
Threshold	—	<5 m	>10 dB	<2 m *	<1.0	<0.85

* Full + Compute SLAM RMSE mean (1.8 m) is within the 2 m budget; maximum across all twenty scenarios is 2.3 m under dense/adverse conditions.

4.1.3. Multi-Objective Communication Cost via Pareto Analysis

The communication cost term J_{comm} (Equation (2)) is extended to a multi-objective formulation incorporating bandwidth, latency, and jitter. An AHP-weighted scalarization is used for real-time MPC integration; the full derivation, pairwise comparison matrix (Supplementary S1), and two-dimensional Pareto-front projections (Supplementary S4) are provided in Supplementary Material S5. Table 11 summarizes the resulting tradeoff surface. The AHP-weighted operating point achieves a balanced Pareto solution: SNR link time decreases by only 0.4 percentage points from the SNR-only baseline while delivering 8.9% higher mean bandwidth, 15.1% lower latency, and 12.9% lower jitter. This tradeoff matters most for collaborative SLAM update windows, where end-to-end latency and jitter directly affect how quickly and reliably landmark observations are exchanged between agents.

Table 11. Pareto front tradeoff across four communication quality metrics. Values are mission averages; improvement is relative to the SNR-only predictive MPC baseline (Table 7). AHP: analytic hierarchy process; SNR: signal-to-noise ratio; BW: available bandwidth; Lat: end-to-end latency.

Operating Point	SNR > 10 dB (%)	Mean BW (Mbps)	Mean Lat (ms)	Jitter (ms)
SNR-only (baseline)	84.1	12.3	38.4	14.7
BW-maximizing	82.6	14.8	36.1	13.2
Latency-minimizing	83.0	13.5	29.8	12.1
AHP-weighted (proposed)	83.7	13.4	32.6	12.8

4.1.4. AHP Weight Sensitivity (Tornado Analysis)

To confirm that the AHP result is not sensitive to the exact weight values, each weight w_i is perturbed by $\pm 20\%$ (keeping the remaining weights proportionally renormalized) and the change in the four output metrics is recorded. Figure 12 shows the resulting tornado chart.




Figure 12. Sensitivity, scalability, and cliff extension analyses. (a) AHP weight sensitivity tornado chart: each weight w_i perturbed by $\pm 20\%$ (remaining weights renormalized); bars show change in composite MPC cost; the largest perturbation (w_{SNR}) shifts cost by only $\pm 3.1\%$, confirming robustness of the AHP-weighted operating point. (b) Empirical QP solve-time scaling of the condensed Riccati recursion on the NVIDIA Jetson Xavier NX, for prediction horizons $N_p \in \{10, \dots, 70\}$ and three state dimensions n_x ; log scale; dashed red line: 100 ms real-time budget; the condensed Riccati solver remains within budget up to $N_p = 280$ (vs. $N_p \approx 40$ for warm-started IPOPT). (c) Effective navigation cliff time t_c^{eff} vs. steady-state good-channel probability π_{good} for the collaborative SLAM fleet ($N = 5$); shaded region: benefit over single-agent baseline; at the validated nominal operating point ($\pi_{good} = 0.789$), t_c^{eff} extends to $\approx 1.6\times$ the baseline 72.9 h cliff. AHP: analytic hierarchy process; QP: quadratic program; IPOPT: interior-point optimizer.

The largest single-weight perturbation ($w_{SNR} \pm 20\%$) shifts the composite cost by only $\pm 3.1\%$ and mean latency by ± 1.8 ms, changes that are operationally negligible given the 50 ms latency threshold. The jitter weight w_{jit} produces the smallest impact ($\pm 0.4\%$),

consistent with its subordinate role in the intelligence, surveillance, and reconnaissance (ISR) priority ordering. These results confirm that the AHP solution is not a fragile optimum dependent on precisely specified weights; any reasonable ISR-domain weight assignment consistent with the SNR-first priority ordering yields a Pareto-optimal operating point within 5% of the reported values.

4.1.5. Structure-Exploiting Condensed Riccati MPC Solver (Implementation Enabler, Not a Contribution)

To meet the 10 Hz real-time requirement of the integrated BAZ II controller on the Jetson Xavier NX, the BAZ IPOPT QP step is replaced by a standard structure-exploiting condensed Riccati factorization drawn from the established sparse-MPC literature [13–15]. Authorship of the Riccati recursion itself is not claimed: the block-tridiagonal factorization is a textbook construction [13] and is also the underlying kernel used by the open-source HPIPM [14] and FORCES [15] solvers. What is reported here is the integration of this factorization as the real-time iteration (RTI) step of the BAZ II MPC, together with the embedded-hardware solve-time profile (Table 12) needed to verify that the four primary contributions C1–C4 fit within the 100 ms Layer 2 budget. Complexity is reduced from $\mathcal{O}(N_p^3)$ to $\mathcal{O}(N_p)$; the implementation details, RTI linearization scheme, and complexity analysis appear in Supplementary Material S5. Table 12 and Figure 12 summarize the performance improvements. Table 12 presents solve-time statistics collected over 10,000 consecutive QP solves during the representative mission scenario on the Jetson Xavier NX. The larger sample size provides reliable characterization of tail probability: the empirical 99th-percentile estimate requires at minimum 1000 samples, and the Gumbel-fitted worst-case at 10^{-4} probability is also reported. IPOPT benchmarks include both cold-started (BAZ condition) and warm-started variants (initialized from the previous solution); warm-starting is standard practice at 10 Hz and reduces IPOPT solve time substantially. The proposed Riccati solver at $N_p = 50$ achieves a mean solve time of 3.1 ms, which is 62% lower than the BAZ IPOPT baseline at the same horizon, and a worst-case of 5.8 ms, well within the 100 ms Layer 2 real-time budget at 10 Hz. This restores the full predictive horizon of BAZ while simultaneously incorporating the additional BAZ II cost terms, removing the $N_p = 20$ compromise imposed by the added GP and SLAM computational load.

Table 12. MPC QP solver comparison on the NVIDIA Jetson Xavier NX (21 TOPS). Statistics collected over 10,000 consecutive QP solves. “WS” denotes warm-started IPOPT (previous-solution initialization); “CS” denotes cold-started. Gumbel WC: Gumbel distribution-fitted worst-case at 10^{-4} exceedance probability. IPOPT: interior point optimizer; Riccati: condensed Riccati recursion (proposed); N_p : prediction horizon; $n_x = 12$, $n_u = 4$.

Solver	N_p	Mean (ms)	99th pct (ms)	Worst (ms)	Gumbel WC (ms)
IPOPT CS (BAZ baseline)	50	8.2	14.2	18.3	21.4
IPOPT WS (warm-started)	50	5.1	9.8	13.2	15.6
IPOPT CS (BAZ II, $N_p = 20$)	20	4.1	7.8	11.2	13.1
Riccati recursion (proposed)	50	3.1	4.9	5.8	6.7

4.1.6. Solve-Time Scaling with Horizon and State Dimension

Figure 12 characterizes the empirical complexity scaling of the Riccati recursion across prediction horizons $N_p \in \{10, \dots, 70\}$ and state dimensions $n_x \in \{6, 12, 18\}$ (10,000 solves each), compared against warm-started IPOPT at the same configurations. The empirical scaling exponents show the theoretically predicted behavior: $\mathcal{O}(N_p)$ for horizon (compared to $\mathcal{O}(N_p^{2.4})$ for IPOPT), and $\mathcal{O}(n_x^3)$ for the number of states (IPOPT has a similar exponent due to the use of a dense Riccati step). The point where the Riccati solver is faster than warm-started IPOPT is at approximately $N_p \approx 8$, showing that the structure advantage comes into effect right away for all practical MPC horizon sizes. The Riccati solver runs in

less than 100 ms real-time until $N_p = 280$ at $n_x = 12$. Therefore, there is considerable room for future increase in the size of the MPC horizon.

5. Discussion

5.1. Significance of the 72-h Navigation Cliff and FDIR

The computational verification of $t_c = 72.9 \pm 2.4$ h corroborates the BAZ analytical result ($t_c = 72.3$ h) nearly exactly. The identified physical mechanism—irreversible accumulation of thermal hysteresis micromechanical creep—is the causal explanation that was absent from the original CTMC characterization.

Beyond confirming the threshold, the mechanism changes how it can be used operationally. This 72-h cliff monitor is a proactive fault detection, isolation, and recovery (FDIR) mechanism that is invoked before INS divergence and spans the physical (thermal) to cyber (communication) domains. Works on LSTM-based real-time planning [95] and cost-effectiveness of autonomous aerial systems [96] give complementary context for the operational deployment of such an approach.

5.2. Collaborative SLAM as a t_c Extension Mechanism

The localization improvement observed for GPS-denied agents illustrates what multi-agent collaboration can offer to long-endurance operations. Its connection to the navigation cliff is direct: by providing external position-aiding that arrests INS estimate divergence, collaborative SLAM turns t_c from a hard failure threshold into a recoverable degradation boundary. The modest drop from full communication (1.07 m) to CTMC-degraded communication (2.06 m) confirms that the information-weighted fusion inherits the same robustness to intermittent link outages that was proved for the scalar CTMC model in BAZ [7].

The sensitivity analyses make this scaling concrete. The fleet size sensitivity curve (Figure 5a) quantifies this scaling: the power-law exponent $\gamma = 0.84$ shows near-linear improvement, and $N \geq 4$ suffices to meet the 2 m error budget under both full and CTMC-degraded communication. The CTMC channel quality heatmap (Figure 5b) identifies the robust operating zone ($\pi_{\text{good}} > 0.65$, RMSE < 2 m), providing channel-quality thresholds for mission planning. The synthesis figure (Section 5.3, Figure 12) directly quantifies the cliff extension: for $N = 5$ and $\pi_{\text{good}} = 0.80$, t_c^{eff} extends beyond the 96 h simulation window, and even the minimum useful configuration ($N = 2$, $\pi_{\text{good}} = 0.55$) yields a 13.5% extension to 82.1 ± 3.8 h [5,97].

5.3. Navigation Cliff Extension as a Function of Channel Quality and Fleet Size

An operational question left open by the individual contribution results is the following: by how much does collaborative SLAM extend the effective navigation cliff t_c , and how does this extension scale with link quality and fleet size? Figure 12 answers this directly. The effective cliff time $t_c^{\text{eff}}(N, \pi_{\text{good}})$ —defined as the flight duration at which position error crosses 200 m (the operational GPS-denied degradation limit)—is computed via Monte Carlo simulation (100 seeds per point) over $\pi_{\text{good}} \in [0.4, 1.0]$ and $N \in \{1, 2, 3, 4, 5\}$. At the validated operating point ($N = 5$, $\pi_{\text{good}} = 0.80$), t_c^{eff} exceeds the 96 h simulation ceiling, indicating that the 72-h cliff is effectively mitigated within the simulated regime by the five-agent fleet under nominal channel conditions.

The extension degrades gracefully as conditions worsen rather than collapsing at a single operating point. Even at $N = 2$, $\pi_{\text{good}} = 0.55$ —the worst-case combination that is still inside the robust SLAM zone of Figure 5b— t_c^{eff} reaches 82.1 ± 3.8 h, a 13.5% extension over the unmitigated BAZ analytical cliff at $t_c = 72.3$ h. These curves serve as a direct mission planning tool: given a target sortie duration and an anticipated channel quality,

a planner can read off the minimum fleet size required to keep the position error within budget for the full mission.

5.4. Predictive Versus Reactive Communication Management

The 1.9 percentage point gain in link availability achieved by the GP-based predictive MPC reflects the compounding benefit described in Section 3.5. Unlike the reactive scalar CTMC model inherited from BAZ, the GP posterior gives the controller spatially resolved foresight, allowing it to steer the vehicle away from high-uncertainty shadowed regions before the instantaneous SNR drops below threshold. This look-ahead capability is critical for protecting the high-SNR windows that collaborative SLAM needs for reliable information exchange—a “communication safety margin” that the original single-agent reactive formulation could not provide by design.

However, this advantage is bounded by the fidelity of the underlying RF map. The predictive gain is largest when the GP posterior is well-conditioned over the planning horizon and shrinks as the vehicle enters previously unsampled terrain, which is precisely where the online recursive update of Section 3.5 becomes important. Therefore, the predictive and reactive modes are complementary rather than mutually exclusive: predictive routing shapes the nominal trajectory, while the reactive penalty remains as a fallback when the map is locally uncertain.

5.5. Anisotropic GP Mapping and Cognitive Radio Resilience

The 48.6% reduction in RF map prediction error obtained by the ARD, spatiotemporal and online GP variant relative to the isotropic baseline shows that directional SNR structure contributes substantially to mapping uncertainty in urban environments. The fitted vertical length scale ($\ell_z = 45$ m) is much shorter than its horizontal counterpart ($\ell_x = 180$ m), reflecting the altitude-dependent LoS probability of the path-loss model (Supplementary Material S5, Section S3.3); an isotropic kernel collapses both scales into a single correlation dimension and incurs the corresponding penalty. The hyperparameter sensitivity heatmap (Section 3.5.3, Figure 8) corroborates this interpretation: the isotropic ridge lies +1.4 dB above the anisotropic optimum, while the optimum itself is robust to $\pm 30\%$ perturbations of (ℓ_x, ℓ_z) (RMSPE rises by only +0.8 dB). The 10 Hz online recursive update is most consequential for extended multi-day deployments, where a statically trained GP would fail to produce accurate maps as the UAV enters previously unsampled terrain [32].

Cognitive-radio performance shows a clear asymmetry between barrage and sweep jamming. Barrage jamming on a single channel is readily mitigated: the 5 ms sensing window is long enough to discriminate the jammed channel from the three clear alternatives, and the mean restoration latency of 340 ms is dominated by the CTMC update from a degraded state estimate to a clear-channel prediction. A detailed JNR sensitivity analysis appears in Section 3.6.4 and Figure 9: link time exceeds 76% up to JNR = 18 dB for barrage jamming and ≈ 14 dB for sweep jamming. Sweep jamming at 2 Hz cycles faster than the sensing window can resolve, restricting the hopping strategy to a reactive rather than a predictive mode. Extending the GP-based hopping policy of Equation (16) with temporal predictions of sweep patterns is a natural direction for future research.

5.6. Multi-Objective Communication Cost and Solver Tractability

The Pareto analysis confirms that the proposed AHP-weighted J_{comm} formulation provides a more operationally balanced link quality metric than the SNR-only BAZ baseline, reducing latency and jitter at negligible SNR cost. A detailed discussion of the multi-objective cost structure and AHP weight sensitivity (Figure 12) is provided in Supplementary Material S5.

The practical value of this reformulation is that latency and jitter, rather than SNR alone, are the quantities that govern the usability of the inter-agent link for collaborative SLAM and control. Treating them as explicit cost axes allows the planner to trade a small SNR margin for a measurable improvement in the metrics that matter operationally. The tornado analysis further shows that this tradeoff is stable under $\pm 20\%$ perturbation of the AHP weights, so the result does not depend on a fragile choice of relative priorities.

5.7. Computational Tractability

The adaptive sparse/full processing mode allows the integrated framework to fit within the computational budget of current embedded hardware. The condensed Riccati solver (Section 4.1.5) reduces worst-case bounded solve times, enabling the combined BAZ II contributions without compromising the MPC horizon length. Detailed empirical complexity scaling results (Figure 12) and real-time computation metrics across processing modes are available in Supplementary Material S5.

A normalized heatmap of all five integration configurations against six performance metrics confirms that the Full Integration configuration achieves the highest aggregate score (82.4%), with the BAZ II contributions acting synergistically rather than competing across the performance space (Supplementary Material S4, Figure S1). This indicates that the integration does not merely sum the individual components but lets them share computational headroom, which is what keeps the combined controller inside the embedded real-time budget.

5.8. Limitations

Several limitations bound the scope of the present results, and are stated here explicitly so that the conclusions can be read in the correct light. They define the regime in which the reported numbers are meaningful and the regime in which they are not.

First, every result reported in this paper is obtained in simulation. No hardware-in-the-loop or field data are presented, and the thermal hysteresis drift model, CTMC channel, and jamming models are all idealizations calibrated against published datasheets rather than measurements from the authors' own platform. Therefore, the quantitative figures should be read as evidence of internal consistency and relative improvement, not as guarantees of field performance.

Second, the Monte Carlo ensembles emphasize mean behavior, and the simulated outcomes are comparatively consistent across seeds. Rare divergent or failed runs, for instance cases in which a loss of SLAM loop closure coincides with a prolonged bad-channel CTMC sojourn, are under-represented at the seed counts used here. Characterizing this tail behavior rather than only the mean is one motivation for the larger seed budgets identified in the axes for future work provided below.

Third, the SLAM benchmark of Section 3.4 is conducted over 30-min visual inertial scenarios, whereas the cliff itself is a multi-day phenomenon. Therefore, the comparison with Swarm-SLAM, Kimera-Multi, and D2SLAM is measured at the benchmark horizon, but is extrapolated at multi-day timescales rather than being measured. In addition, those systems are not designed for the stochastic CTMC degradation regime, so the comparison characterizes their behavior under a regime outside their design intent rather than providing a like-for-like ranking.

Finally, the computational feasibility claim is established for a single embedded platform (the NVIDIA Jetson Xavier NX) and a fleet of five agents. Scaling to larger swarms or to different hardware, along with the decentralized consensus and compression that this would require, is left to future work.

6. Conclusions

This paper presents BAZ II as a simulation-validated multi-agent extension of the BAZ analytical framework. The unifying idea of the paper is that the 72-h navigation cliff need not be treated as a fixed endurance limit; when communication quality is promoted to a first-class planning resource and combined with multi-agent collaboration, the cliff can instead be reframed as a degradation event that mission planners can anticipate and manage.

Four ideas give this reframing substance. A physically grounded thermal-hysteresis drift model explains why the cliff exists, rather than merely predicting that it does. A communication-aware collaborative SLAM filter lets agents aid one another's localization through an explicitly stochastic channel. An RF-aware predictive controller supported by a cognitive radio layer gives the fleet spatial and spectral foresight in place of purely reactive recovery. Finally, a multi-objective communication cost replaces a single signal-strength proxy with the link qualities that actually govern cooperative navigation. The contribution of the paper is not any one of these components in isolation but their integration into a single hierarchical system that remains within the real-time budget of embedded UAV hardware.

The central claim that collaboration and proactive communication management can extend the effective navigation cliff is supported consistently across the simulation study; however, it remains a simulation result. Translating it into an operational guarantee will require the hardware-in-the-loop and long-endurance field testing set out in Section 5.8 together with multi-day SLAM characterization, explicit modeling of accelerometer thermal bias, and scaling to larger and heterogeneous fleets. Pursued along these axes, BAZ II is best understood as a design baseline for long-endurance GPS-denied multi-UAV navigation rather than a finished operational solution. The associated key performance indicator and optimization methodology builds upon the authors' prior work in aerospace systems engineering [98,99].

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/drones10080561/s1>. S1: AHP Pairwise Comparison Matrix for MPC Cost Weights; S2: Full CTMC Parameter Derivation and Lumpability Proof; S3: Representative Simulation Run Scripts; S4: Two-Dimensional Pareto-Front Projections; S5: Extended Analyses and Derivations.

Author Contributions: Conceptualization, M.A. (Mosab Alrashed); methodology, M.A. (Mosab Alrashed); software, M.A. (Mosab Alrashed); formal analysis, M.A. (Mosab Alrashed); validation, M.A. (Mosab Alrashed); investigation, M.A. (Mosab Alrashed) and H.A.; resources, H.A. and M.A. (Mohammad Alqattan); data curation, M.A. (Mosab Alrashed); writing—original draft preparation, M.A. (Mosab Alrashed); writing—review and editing, H.A. and M.A. (Mohammad Alqattan); visualization, M.A. (Mosab Alrashed); supervision, M.A. (Mohammad Alqattan); project administration, M.A. (Mosab Alrashed) and M.A. (Mohammad Alqattan); funding acquisition, M.A. (Mohammad Alqattan). All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Kuwait Foundation for the Advancement of Science (KFAS) under Application Number 2307. The article processing charge (APC) was funded by the American International University (AIU) institutional open-access publishing fund.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: All data generated or analyzed during this study are included in this published article and its Supplementary Information Files. The precursor BAZ study is published as Alrashed et al. [7] (*Drones* 2026, 10(5), 371).

Acknowledgments: The authors thank the anonymous reviewers of BAZ [7] whose constructive feedback directly motivated the four primary contributions of this work, and the international

UAV research community for shared software libraries and documentation that accelerated this program. The authors also thank the four anonymous reviewers of the present paper, whose detailed second-round critiques led directly to the consolidation of the contribution list, the addition of the quantitative SLAM benchmark (Table 5), the regeneration of Figures 8 and 9, and the clarification of the BAZ acronym, the MPC framework justification, and the accelerometer thermal error budget.

Conflicts of Interest: The authors declare no conflicts of interest; the funders had no role in the design of the study, in the collection, analysis, or interpretation of data, in the writing of the manuscript, or in the decision to publish the results.

Nomenclature

A complete symbol and parameter table is provided in Supplementary Material S5 (Section S5.1). The principal symbols and abbreviations used throughout this paper are listed below

Latin Symbols

t_c	Navigation cliff threshold (h)	N_p	MPC prediction horizon
t_c^{eff}	Effective cliff threshold with SLAM (h)	N	Fleet size (number of agents)
$\dot{\theta}$	Gyroscope drift rate	J	Unified MPC cost
b_0	Turn-on bias	Q, R	State/control weight matrices
b_{rw}	Gyro random walk	G	CTMC generator matrix
b_{th}	Thermal sensitivity bias	π_{good}	Steady-state good-channel availability
b_{hyst}	Thermal hysteresis bias	$\lambda_{gb}, \lambda_{bg}$	CTMC transition rates
α	Bifurcation amplification coefficient	w_c, w_i, w_κ	MPC sub-cost weights
β^*	Bifurcation threshold	ℓ_x, ℓ_z	GP horizontal/vertical length scales
κ_T	Thermal sensitivity coefficient	σ^2	GP posterior variance
κ_H	Hysteresis accumulation coefficient	μ_{GP}	GP posterior mean SNR
T_0, T_A, T_P	Temperature offset, amplitude, period	η_{ED}	Energy-detector decision threshold
ε	MPC sub-optimality bound	P_{FA}, P_D	False-alarm and detection probabilities
λ	Lyapunov stability rate	c_{total}	Normalized computational load
x, u	State and control vectors	δ_{trig}	Event-trigger displacement threshold

Abbreviations

AHP	Analytic hierarchy process	LoS/NLoS	Line-of-sight/non-line-of-sight
ARD	Automatic relevance determination	MEMS	Microelectromechanical systems
CEP	Circular error probable	MPC	Model predictive control
CR	Cognitive radio	ODE	Ordinary differential equation
CTMC	Continuous-time Markov chain	QP	Quadratic program
CPU	Central processing unit	RF	Radio frequency
EKF	Extended Kalman filter	RMS/RMSE	Root mean square (error)
FDIR	Fault detection, isolation, and recovery	RMSPE	Root mean square prediction error
GP	Gaussian process	RTI	Real-time iteration
GPS	Global Positioning System	SINR	Signal-to-interference-plus-noise ratio
INS	Inertial navigation system	SLAM	Simultaneous localization and mapping
IQR	Interquartile range	SNR	Signal-to-noise ratio
IPOPT	Interior-point optimizer	TOPS	Tera operations per second
ISR	Intelligence, surveillance, and reconnaissance	UAV	Unmanned aerial vehicle
JNR	Jammer-to-noise ratio		

References

1. Airbus Defence and Space. Zephyr High Altitude Platform Station (HAPS). Manufacturer Official Website. 2023. Available online: <https://www.airbus.com/en/products-services/defence/uas/zephyr> (accessed on 8 December 2025).
2. Tewari, A. *Atmospheric and Space Flight Dynamics: Modeling and Simulation with MATLAB and Simulink*; Birkhäuser: Boston, MA, USA, 2007.
3. Beard, R.W.; McLain, T.W. *Small Unmanned Aircraft: Theory and Practice*; Princeton University Press: Princeton, NJ, USA, 2012. [CrossRef]

4. Oettershagen, P.; Melzer, A.; Mantel, T.; Rudin, K.; Lotz, R.; Siebenmann, D.; Leutenegger, S.; Alexis, K.; Siegwart, R. A Solar-Powered Hand-Launchable UAV for Low-Altitude Multi-Day Continuous Flight. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA, 26–30 May 2015; pp. 3986–3993. [\[CrossRef\]](#)
5. Oettershagen, P.; Melzer, A.; Mantel, T.; Rudin, K.; Stastny, T.; Wawrzacz, B.; Hinzmann, T.; Leutenegger, S.; Alexis, K.; Siegwart, R. Design of Small Hand-Launched Solar-Powered UAVs: From Concept Study to a Multi-Day World Endurance Record Flight. *J. Field Robot.* **2017**, *34*, 1352–1377. [\[CrossRef\]](#)
6. Chu, Y.; Ho, C.; Lee, Y.; Li, B. Development of a Solar-Powered Unmanned Aerial Vehicle for Extended Flight Endurance. *Drones* **2021**, *5*, 44. [\[CrossRef\]](#)
7. Alrashed, M.; Fenjan, A.; Aldaihani, H.; Alqattan, M. Stochastically Optimal Hierarchical Control for Long-Endurance UAVs Under Communication Degradation: Theory and Validation. *Drones* **2026**, *10*, 371. [\[CrossRef\]](#)
8. Guo, G.; Chai, B.; Cheng, R.; Wang, Y. Temperature Drift Compensation of a MEMS Accelerometer Based on DLSTM and ISSA. *Sensors* **2023**, *23*, 1809. [\[CrossRef\]](#) [\[PubMed\]](#)
9. El-Sheimy, N.; Youssef, A. Inertial Sensors Technologies for Navigation Applications: State of the Art and Future Trends. *Satell. Navig.* **2020**, *1*, 2. [\[CrossRef\]](#)
10. Analog, D. ADIS16488: Tactical Grade Ten Degree of Freedom Inertial Sensor; Analog Devices Inc: Norwood, MA, USA, 2022.
11. Yang, G.; Xiao, M.; Peng, Y.; Ma, Z.; Wang, C. Enhanced Neural Network-Based Thermal Compensation Method for MEMS IMUs. In *Proceedings of the 2023 6th International Conference on Mechatronics, Robotics and Automation (ICMRA)*, Xiamen, China, 17–19 November 2023; pp. 184–189.
12. Snelson, E.; Ghahramani, Z. Sparse Gaussian Processes Using Pseudo-inputs. In *Proceedings of the Advances in Neural Information Processing Systems*; MIT Press: Cambridge, MA, USA, 2006; Volume 18, pp. 1257–1264.
13. Rawlings, J.B.; Mayne, D.Q.; Diehl, M.M. *Model Predictive Control: Theory, Computation, and Design*, 2nd ed.; Nob Hill Publishing: Madison, WI, USA, 2017.
14. Frison, G.; Diehl, M. HPIPM: A High-Performance Quadratic Programming Framework for Model Predictive Control. *IFAC-PapersOnLine* **2020**, *53*, 6563–6569. [\[CrossRef\]](#)
15. Domahidi, A.; Zraggen, A.U.; Zeilinger, M.N.; Morari, M.; Jones, C.N. Efficient interior point methods for multistage problems arising in receding horizon control. In *Proceedings of the 2012 IEEE 51st Conference on Decision and Control (CDC)*, Maui, HI, USA, 10–13 December 2012; pp. 668–674. [\[CrossRef\]](#)
16. Liberzon, D. *Switching in Systems and Control*; Birkhäuser: Boston, MA, USA, 2003.
17. Shorten, R.; Wirth, F.; Mason, O.; Wulff, K.; King, C. Stability Criteria for Switched and Hybrid Systems. *SIAM Rev.* **2007**, *49*, 545–592. [\[CrossRef\]](#)
18. Khalil, H.K. *Nonlinear Systems*, 3rd ed.; Prentice Hall: Upper Saddle River, NJ, USA, 2002.
19. Wen, J.T.-Y.; Kreutz-Delgado, K. The Attitude Control Problem. *IEEE Trans. Autom. Control* **1991**, *36*, 1148–1162. [\[CrossRef\]](#)
20. Xu, H.; Liu, P.; Chen, X.; Shen, S. D²SLAM: Decentralized and Distributed Collaborative Visual-Inertial SLAM System for Aerial Swarm. *IEEE Trans. Robot.* **2024**, *40*, 3445–3464. [\[CrossRef\]](#)
21. Lajoie, P.Y.; Beltrame, G. Swarm-SLAM: Sparse Decentralized Collaborative Simultaneous Localization and Mapping Framework for Multi-Robot Systems. *IEEE Robot. Autom. Lett.* **2024**, *9*, 475–482. [\[CrossRef\]](#)
22. Tian, Y.; Chang, Y.; Herrera Arias, F.; Nieto-Granda, C.; How, J.P.; Carlone, L. Kimera-Multi: Robust, Distributed, Dense Metric-Semantic SLAM for Multi-Robot Systems. *IEEE Trans. Robot.* **2022**, *38*, 2022–2038. [\[CrossRef\]](#)
23. Lajoie, P.Y.; Ramtoula, B.; Chang, Y.; Carlone, L.; Beltrame, G. DOOR-SLAM: Distributed, Online, and Outlier Resilient SLAM for Robotic Teams. *IEEE Robot. Autom. Lett.* **2020**, *5*, 1656–1663. [\[CrossRef\]](#)
24. Wu, Q.; Zeng, Y.; Zhang, R. Joint Trajectory and Communication Design for Multi-UAV Enabled Wireless Networks. *IEEE Trans. Wirel. Commun.* **2018**, *17*, 2109–2121. [\[CrossRef\]](#)
25. Chen, X.; Zhong, X.; Zhang, Z.; Dai, L.; Zhou, S. High-Efficiency Urban 3D Radio Map Estimation Based on Sparse Measurements. *IEEE Trans. Veh. Technol.* **2025**, *74*, 16488–16493. [\[CrossRef\]](#)
26. Eskandari, M.; Huang, H.; Savkin, A.V.; Ni, W. Model Predictive Control-Based 3D Navigation of a RIS-Equipped UAV for Line-of-Sight Wireless Communication with a Ground Intelligent Vehicle. *IEEE Trans. Intell. Veh.* **2023**, *8*, 2371–2384. [\[CrossRef\]](#)
27. Haykin, S. Cognitive Radio: Brain-Empowered Wireless Communications. *IEEE J. Sel. Areas Commun.* **2005**, *23*, 201–220. [\[CrossRef\]](#)
28. Mitola, J.; Maguire, G.Q. Cognitive Radio: Making Software Radios More Personal. *IEEE Pers. Commun.* **1999**, *6*, 13–18. [\[CrossRef\]](#)
29. Zhao, Q.; Sadler, B.M. A Survey of Dynamic Spectrum Access. *IEEE Signal Process. Mag.* **2007**, *24*, 79–89. [\[CrossRef\]](#)
30. Miettinen, K. *Nonlinear Multiobjective Optimization*; Springer: New York, NY, USA, 1999. [\[CrossRef\]](#)
31. Saaty, T.L. *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*; McGraw-Hill: New York, NY, USA, 1980.
32. Rasmussen, C.E.; Williams, C.K.I. *Gaussian Processes for Machine Learning*; MIT Press: Cambridge, MA, USA, 2006.
33. Titterton, D.H.; Weston, J.L. *Strapdown Inertial Navigation Technology*, 2nd ed.; Institution of Engineering and Technology: Stevenage, UK, 2004.

34. Groves, P.D. *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*, 2nd ed.; Artech House: Norwood, MA, USA, 2013.
35. Cui, R.; Xu, J.; Huang, B.; Xu, H.; Peng, M.; Yang, J.; Zhang, J.; Gu, Y.; Chen, D.; Li, H.; et al. A Temperature Compensation Approach for Micro-Electro-Mechanical Systems Accelerometer Based on Gated Recurrent Unit–Attention and Robust Local Mean Decomposition–Sample Entropy–Time-Frequency Peak Filtering. *Micromachines* **2024**, *15*, 483. [\[CrossRef\]](#) [\[PubMed\]](#)
36. Zhuo, C.; Du, J.; Tang, H.; Liu, Q. Special Thermal Compensation Experiment and Algorithm Design for Inertial Navigation System. In *Proceedings of the 2019 DGON Inertial Sensors and Systems (ISS)*; IEEE: Braunschweig, Germany, 2019; pp. 1–16. [\[CrossRef\]](#)
37. Shuster, M.D. A Survey of Attitude Representations. *J. Astronaut. Sci.* **1993**, *41*, 439–517.
38. Solà, J. Quaternion Kinematics for the Error-State Kalman Filter. *arXiv* **2017**, arXiv:1711.02508.
39. Lateef, F.; Kas, M.; Ruichek, Y. From GPS to AI: A Comprehensive Review of Unmanned Aerial Vehicle (UAV) Localization Solutions. *ISPRS J. Photogramm. Remote Sens.* **2025**, *230*, 402–451. [\[CrossRef\]](#)
40. Lu, Y.; Xue, Z.; Xia, G.-S.; Zhang, L. A Survey on Vision-Based UAV Navigation. *Geo-Spat. Inf. Sci.* **2018**, *21*, 21–32. [\[CrossRef\]](#)
41. Chang, Y.; Cheng, Y.; Manzoor, U.; Murray, J. A Review of UAV Autonomous Navigation in GPS-Denied Environments. *Robot. Auton. Syst.* **2023**, *170*, 104533. [\[CrossRef\]](#)
42. Vanegas, F.; Gonzalez, F. Enabling UAV Navigation with Sensor and Environmental Uncertainty in Cluttered and GPS-Denied Environments. *Sensors* **2016**, *16*, 666. [\[CrossRef\]](#) [\[PubMed\]](#)
43. Al-Jarrah, O.Y.; Shatnawi, A.S.; Shurman, M.M.; Ramadan, O.A.; Muhaidat, S. Exploring Deep Learning-Based Visual Localization Techniques for UAVs in GPS-Denied Environments. *IEEE Access* **2024**, *12*, 113049–113071. [\[CrossRef\]](#)
44. Gao, Y.; Wang, Y.; Tian, L.; Li, D.; Wang, F. Visual Navigation Algorithms for Aircraft Fusing Neural Networks in Denial Environments. *Sensors* **2024**, *24*, 4797. [\[CrossRef\]](#) [\[PubMed\]](#)
45. UAVs@Berkeley Research Group. GPS-Denied Navigation Using Visual SLAM for Quadrotor UAVs. Project Technical Documentation. 2024. Available online: <https://uav.studentorg.berkeley.edu/106a-fa24-gps/> (accessed on 8 December 2025).
46. Niculescu, V.; Polonelli, T.; Magno, M.; Benini, L. NanoSLAM: Enabling Fully Onboard SLAM for Tiny Robots. *IEEE Internet Things J.* **2024**, *11*, 13584–13607. [\[CrossRef\]](#)
47. Teague, S.; Chahl, J. An Algorithm for Affordable Vision-Based GNSS-Denied Strapdown Celestial Navigation. *Drones* **2024**, *8*, 652. [\[CrossRef\]](#)
48. Xin, L.; Tang, Z.; Gai, W.; Liu, H. Vision-Based Autonomous Landing for the UAV: A Review. *Aerospace* **2022**, *9*, 634. [\[CrossRef\]](#)
49. Lippiello, V.; Cacace, J. Robust Visual Localization of a UAV Over a Pipe-Rack Based on the Lie Group SE(3). *IEEE Robot. Autom. Lett.* **2022**, *7*, 295–302. [\[CrossRef\]](#)
50. Burri, M.; Nikolic, J.; Gohl, P.; Schneider, T.; Rehder, J.; Omari, S.; Achtelik, M.W.; Siegwart, R. The EuRoC Micro Aerial Vehicle Datasets. *Int. J. Robot. Res.* **2016**, *35*, 1157–1163. [\[CrossRef\]](#)
51. Robicquet, A.; Sadeghian, A.; Alahi, A.; Savarese, S. Learning Social Etiquette: Human Trajectory Understanding in Crowded Scenes. In *Proceedings of the European Conference on Computer Vision (ECCV)*; Springer Nature: Amsterdam, The Netherlands, 2016; pp. 549–565. [\[CrossRef\]](#)
52. Cadena, C.; Carlone, L.; Carrillo, H.; Latif, Y.; Scaramuzza, D.; Neira, J.; Reid, I.; Leonard, J.J. Past, Present, and Future of Simultaneous Localization and Mapping: Toward the Robust-Perception Age. *IEEE Trans. Robot.* **2016**, *32*, 1309–1332. [\[CrossRef\]](#)
53. Kelly, D.; Majda, A.J.; Tong, X.T. Concrete Ensemble Kalman Filters with Rigorous Catastrophic Filter Divergence. *Proc. Natl. Acad. Sci. USA* **2015**, *112*, 10589–10594. [\[CrossRef\]](#) [\[PubMed\]](#)
54. Specht, M. A Statistical Distribution Analysis of Navigation Positioning System Errors—Issue of the Empirical Sample Size. *Sensors* **2020**, *20*, 7144. [\[CrossRef\]](#) [\[PubMed\]](#)
55. Tian, Y.; Khosoussi, K.; Rosen, D.M.; How, J.P. Distributed Certifiably Correct Pose-Graph Optimization. *IEEE Trans. Robot.* **2021**, *37*, 2137–2156. [\[CrossRef\]](#) [\[PubMed\]](#)
56. Yao, F.; Liu, Y.; Zhang, W.; Zhu, Z.; Li, C.; Liu, N.; Hu, P.; Yue, Y.; Wei, K.; He, X.; et al. AeroVerse-Review: Comprehensive Survey on Aerial Embodied Vision-and-Language Navigation. *Innov. Inf.* **2025**, *1*, 100015. [\[CrossRef\]](#)
57. Abujoub, S.; McPhee, J.; Irani, R.A. Methodologies for Landing Autonomous Aerial Vehicles on Maritime Vessels. *Aerosp. Sci. Technol.* **2020**, *106*, 106169. [\[CrossRef\]](#)
58. Zhang, L.; Zhang, R. Radio Map-Based 3D Path Planning for Cellular-Connected UAV. *IEEE Trans. Wirel. Commun.* **2021**, *20*, 1975–1989. [\[CrossRef\]](#)
59. Mozaffari, M.; Saad, W.; Bennis, M.; Nam, Y.H.; Debbah, M. A Tutorial on UAVs for Wireless Networks: Applications, Challenges, and Open Problems. *IEEE Commun. Surv. Tutor.* **2019**, *21*, 2334–2360. [\[CrossRef\]](#)
60. Al-Hourani, A.; Kandeepan, S.; Lardner, S. Optimal LAP Altitude for Maximum Coverage. *IEEE Wirel. Commun. Lett.* **2014**, *3*, 569–572. [\[CrossRef\]](#)
61. Rappaport, T.S. *Wireless Communications: Principles and Practice*, 2nd ed.; Prentice Hall: Upper Saddle River, NJ, USA, 2002.

62. Zeng, Y.; Xu, J.; Zhang, R. Energy Minimization for Wireless Communication With Rotary-Wing UAV. *IEEE Trans. Wirel. Commun.* **2019**, *18*, 2329–2345. [[CrossRef](#)]
63. Liu, X.; Lai, B.; Lin, B.; Leung, V.C.M. Joint Communication and Trajectory Optimization for Multi-UAV Enabled Mobile Internet of Vehicles. *IEEE Trans. Intell. Transp. Syst.* **2022**, *23*, 15354–15366. [[CrossRef](#)]
64. Lyu, Y.; Wang, W. Modeling Time-Variant Multipath Number for Geometry-Based Stochastic UAV-to-Ground Channel Model in a Typical Rural Area. *IEEE Antennas Wirel. Propag. Lett.* **2026**, *25*, 1289–1293. [[CrossRef](#)]
65. Wang, H.S.; Moayeri, N. Finite-State Markov Channel—A Useful Model for Radio Communication Channels. *IEEE Trans. Veh. Technol.* **1995**, *44*, 163–171. [[CrossRef](#)]
66. Xu, C.; Zhang, K.; Jiang, Y.; Niu, S.; Yang, T.; Song, H. Communication Aware UAV Swarm Surveillance Based on Hierarchical Architecture. *Drones* **2021**, *5*, 33. [[CrossRef](#)]
67. Wang, Z.; Han, Z.; Tayyaba, S. Adaptive Control for Uncrewed Aerial Vehicles Based on Communication Information Optimization in Complex Environments. *PeerJ Comput. Sci.* **2024**, *10*, e1920. [[CrossRef](#)] [[PubMed](#)]
68. Zhang, F.; Zhou, M.; Qi, L. Distributed and Coordinated Model Predictive Control for Channel Resource Allocation in Cooperative Vehicle Safety Systems. *IEEE Internet Things J.* **2024**, *11*, 19328–19343. [[CrossRef](#)]
69. Yu, Y.; Zheng, X.; He, S. Communication-Aware Asynchronous Distributed Trajectory Optimization for UAV Swarm. *arXiv* **2025**, arXiv:2511.14994. [[CrossRef](#)]
70. Cook, M.V. *Flight Dynamics Principles*, 3rd ed.; Butterworth-Heinemann: Oxford, UK, 2012.
71. Stevens, B.L.; Lewis, F.L.; Johnson, E.N. *Aircraft Control and Simulation: Dynamics, Controls Design, and Autonomous Systems*, 3rd ed.; John Wiley & Sons: Hoboken, NJ, USA, 2015.
72. Lee, T.; Leok, M.; McClamroch, N.H. Geometric Tracking Control of a Quadrotor UAV on SE(3). In *Proceedings of the IEEE Conference on Decision and Control (CDC)*; IEEE: New York, NY, USA, 2010; pp. 5420–5425. [[CrossRef](#)]
73. Bullo, F.; Lewis, A.D. *Geometric Control of Mechanical Systems*; Springer Science & Business Media: New York, NY, USA, 2004.
74. U.S. Department of Defense. *Military Standard: Flying Qualities of Piloted Aircraft*; Technical Report MIL-STD-1797A; U.S. Department of Defense: Washington, DC, USA, 1990.
75. U.S. Department of Defense. *Flying Qualities of Piloted Airplanes*; Technical Report MIL-F-8785C; U.S. Department of Defense: Washington, DC, USA, 1980.
76. SBG Systems. UAV & Unmanned Aerial Vehicles Navigation–Defense. Product Catalog. 2024. Available online: <https://www.sbg-systems.com/defense/uav-unmanned-aerial-vehicles-navigation-defense/> (accessed on 8 December 2025).
77. DJI Enterprise. DJI Matrice 400-Specs. DJI Enterprise Official Documentation. 2024. Available online: <https://enterprise.dji.com/matrice-400/specs> (accessed on 25 December 2025).
78. NovAtel Inc. *NovAtel's SPAN® GNSS/INS with the Analog Devices ADIS16488 IMU Performance Analysis*; Technical Report; NovAtel Inc.: Calgary, AB, Canada, 2012.
79. Xu, Y.; Liu, S.; He, C.; Wu, H.; Cheng, L.; Yan, G.; Huang, Q. Reliability of MEMS Inertial Devices in Mechanical and Thermal Environments: A Review. *Heliyon* **2024**, *10*, e27481. [[CrossRef](#)] [[PubMed](#)]
80. Li, Z.; Cui, Y.; Gu, Y.; Wang, G.; Yang, J.; Chen, K.; Cao, H. Temperature Drift Compensation for Four-Mass Vibration MEMS Gyroscope Based on EMD and Hybrid Filtering Fusion Method. *Micromachines* **2023**, *14*, 971. [[CrossRef](#)] [[PubMed](#)]
81. Wang, Z.; Cheng, X.; Du, J. Thermal Modeling and Calibration Method in Complex Temperature Field for Single-Axis Rotational Inertial Navigation System. *Sensors* **2020**, *20*, 384. [[CrossRef](#)] [[PubMed](#)]
82. Xiong, Z.; Wei, G.; Gao, C.; Long, X. Precision Temperature Control for the Laser Gyro Inertial Navigation System in Long-Endurance Marine Navigation. *Sensors* **2021**, *21*, 4119. [[CrossRef](#)] [[PubMed](#)]
83. Wann, M.; Yen, D.; Gold, H.J. Evaluation and Calibration of Three Models for Daily Cycle of Air Temperature. *Agric. For. Meteorol.* **1985**, *34*, 121–128. [[CrossRef](#)]
84. Lu, L.; Zhou, X.M. A Four-Parameter Model for Estimating Diurnal Temperature Cycle from MODIS Land Surface Temperature Product. *J. Geophys. Res. Atmos.* **2021**, *126*, e2020JD033855. [[CrossRef](#)]
85. Metropolis, N.; Ulam, S. The Monte Carlo Method. *J. Am. Stat. Assoc.* **1949**, *44*, 335–341. [[CrossRef](#)] [[PubMed](#)]
86. Robert, C.P.; Casella, G. *Monte Carlo Statistical Methods*, 2nd ed.; Springer: New York, NY, USA, 2004.
87. Hammersley, J.M.; Handscomb, D.C. *Monte Carlo Methods*; Methuen: London, UK, 1964.
88. Alrashed, M.; Nikolaidis, T.; Pilidis, P.; Jafari, S. Turboelectric Uncertainty Quantification and Error Estimation in Numerical Modelling. *Appl. Sci.* **2020**, *10*, 1805. [[CrossRef](#)]
89. The MathWorks, Inc. MATLAB and Simulink, Release 2023b. 2023. Available online: <https://www.mathworks.com> (accessed on 8 December 2025).
90. Kemeny, J.G.; Snell, J.L. *Finite Markov Chains*; Springer: New York, NY, USA, 1976. [[CrossRef](#)]
91. Costa, O.L.V.; Fragoso, M.D.; Todorov, M.G. *Continuous-Time Markov Jump Linear Systems*; Springer: Berlin, Germany, 2013.
92. Heemels, W.P.M.H.; Johansson, K.H.; Tabuada, P. An Introduction to Event-Triggered and Self-Triggered Control. In *Proceedings of the 2012 IEEE 51st Conference on Decision and Control (CDC)*, Maui, HI, USA, 10–13 December 2012; pp. 3270–3285. [[CrossRef](#)]

93. Cassandras, C.G.; Lafortune, S. *Introduction to Discrete Event Systems*, 2nd ed.; Springer Science & Business Media: New York, NY, USA, 2008.
94. Cieslewski, T.; Choudhary, S.; Scaramuzza, D. Data-Efficient Decentralized Visual SLAM. In Proceedings of the 2018 IEEE International Conference on Robotics and Automation (ICRA), Brisbane, Australia, 21–25 May 2018; pp. 2466–2473. [[CrossRef](#)]
95. Zhang, J.; Guo, Y.; Zheng, L.; Yang, Q.; Shi, G.; Wu, Y. Real-Time UAV Path Planning Based on LSTM Network. *J. Syst. Eng. Electron.* **2024**, *35*, 374–385. [[CrossRef](#)]
96. Everly, R.E.; Limmer, D.C.; MacKenzie, C.A. *Cost-Effectiveness Analysis of Autonomous Aerial Platforms and Communication Payloads*; Technical Report; Iowa State University, Industrial and Manufacturing Systems Engineering: Ames, IA, USA, 2015.
97. Yang, Y.; Khalife, J.; Morales, J.J.; Kassas, Z.M. UAV Waypoint Opportunistic Navigation in GNSS-Denied Environments. *IEEE Trans. Aerosp. Electron. Syst.* **2022**, *58*, 663–678. [[CrossRef](#)]
98. Alrashed, M.; Nikolaidis, T.; Pilidis, P.; Jafari, S.; Alrashed, W. Key Performance Indicators for Turboelectric Distributed Propulsion. *Int. J. Product. Perform. Manag.* **2022**, *71*, 1989–2008. [[CrossRef](#)]
99. Musa, G.; Igie, U.; Di Lorenzo, G.; Alrashed, M.; Navaratne, R. Gas Turbine Compressor Washing Economics and Optimization Using Genetic Algorithm. *J. Eng. Gas Turbines Power* **2022**, *144*, 091012. [[CrossRef](#)]

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.