Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Role Partitioned Swarm Intelligence for Resilient UAV Coordination in Contested Electromagnetic Environments: A Soccer Field Swarm Formation Framework

Mohamed Hijazy Shazin Hassan
shazhijazy@gmail.com
Andhra University
Bobbadi Jaswanth Kumar
Andhra University
Neeraj Sree Mailee Dudaboyina
Andhra University
Giri Raja Sekhar Gunnu
Andhra University
Kuda Nageswara Rao
Andhra University

Research Article

Keywords: UAV swarms; swarm intelligence; role-partitioned coordination; distributed control; electronic warfare resilience; control barrier functions; distributed estimation; software-in-the-loop

Posted Date: November 6th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11485152/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Role Partitioned Swarm Intelligence for Resilient UAV Coordination in Contested Electromagnetic Environments: A Soccer Field Swarm Formation Framework

Mohamed Hijazy Shazin Hassan[1,3,*], Bobbadi Jaswanth Kumar[1], Neeraj Sree Mailee Dudaboyina[1], Giri Raja Sekhar Gunnu[2,3], and Kuda Nageswara Rao[1]

[1]Department of Computer Science & Systems Engineering, Andhra University, Visakhapatnam 530003, Andhra Pradesh, India

[2]Department of Marine Engineering, Andhra University, Visakhapatnam 530003, Andhra Pradesh, India

[3]Vaishvik Vertex Solutions Pvt. Ltd., Srikakulam 532421, Andhra Pradesh, India

*Corresponding author: Mohamed Hijazy Shazin Hassan; email: shazhijazy@gmail.com

Abstract

Autonomous unmanned aerial vehicle (UAV) swarms must coordinate asset protection, threat interception, and collision avoidance even when electronic warfare degrades inter-agent communication. This paper presents Role-Partitioned Swarm Intelligence (RPSI) and its Soccer Field Swarm Formation (SFSF) implementation, which organise agents into Defender, Midfielder, and Forward roles for asset shielding, interception and relay, and outer-layer engagement. The proposed architecture combines rigid-body dynamics on SE(3), Covariance Intersection (CI) for distributed estimation, role-conditioned artificial potential fields, a decentralised High-Order Control Barrier Function (HO-CBF) quadratic-program safety filter, and a communication-health-triggered resilience mode that freezes role assignments and continues with local sensing. A dwell-time-constrained switched-system analysis is provided under stated feasibility and bounded-switching assumptions. In the PX4-based Software-in-the-Loop (SITL) evaluation, RPSI-SFSF achieved 92.7% asset-protection effectiveness (APE) under nominal communication, compared with 86.1% for distributed model predictive control, and 79.8% versus 57.6% under the C5 condition with 0% packet availability. In the 120-run reduced-order packet-drop experiment, RPSI exceeded the homogeneous controller by 8.75, 3.75, and 23.75 percentage points at 100%, 50%, and 0% target-state update packet delivery, respectively. Across-run standard deviations ranged from 10.26 to 20.32 points, and some runs at 100% and 50% delivery fell below the 5.0 m design threshold.

Keywords: UAV swarms; swarm intelligence; role-partitioned coordination; distributed control; electronic warfare resilience; control barrier functions; distributed estimation; software-in-the-loop

1. Introduction

1.1 Motivation

Coordinated swarms of unmanned aerial vehicles (UAVs) have progressed from a theoretical construct to a practical technology used in surveillance, search and rescue, disaster response, environmental monitoring, infrastructure inspection, and coordinated defence. Relative to single-platform systems, swarms offer redundancy, scalability, distributed sensing, and fault tolerance. Realising these benefits, however, requires solving a multi-objective coordination problem involving formation maintenance, collision avoidance, target tracking, resource allocation, communication management, and mission execution under environmental uncertainty.

The difficulty is compounded in contested electromagnetic environments. Electronic warfare (EW) systems degrade or deny communication through jamming, spoofing, packet erasure, and network fragmentation. Coordination architectures that assume persistent connectivity, particularly consensus-based schemes that require repeated state exchange to maintain convergence, experience accumulating estimation error and degraded coordination as link quality deteriorates. Centralised architectures additionally expose single points of failure.

1.2 Limitations of homogeneous coordination

A structural assumption shared by most existing swarm-control paradigms is functional homogeneity: all agents execute identical coordination policies while pursuing multiple mission objectives simultaneously. Each agent must therefore balance asset protection, threat interception, formation maintenance, relaying, and engagement within a single controller. As mission complexity grows, this produces three problems:

1. **Resource-allocation conflict.** Simultaneously active objective gradients can interfere; Section 3.4 formalises this as

objective interference and shows that removing conflicting pairs cannot decrease efficiency under a common action budget.
2. **Spatial inefficiency.** Agents pursuing a common blended objective tend to cluster or distribute sub-optimally across defensive, transitional, and offensive regions of the battlespace.
3. **Connectivity dependence.** Homogeneous distributed methods typically rely on frequent state exchange; when communication is denied, every agent's policy degrades because every policy depends on shared state.

This paper addresses the question: *can mission effectiveness be improved by structurally partitioning a swarm into specialised functional groups rather than coordinating all agents as a homogeneous collective?*

1.3 Contribution: Role-Partitioned Swarm Intelligence

We answer this question affirmatively, with qualifications established by experiment, through Role-Partitioned Swarm Intelligence (RPSI). RPSI treats functional specialisation as a primary organisational principle governing swarm structure, communication patterns, and decision-making, rather than as a secondary task-allocation mechanism. Its concrete implementation, the Soccer Field Swarm Formation (SFSF) framework, partitions the swarm into three role classes inspired by the layered organisation of association football:

- **Defenders** protect high-value assets through asset-centred potentials and screening points located on the asset–threat segment;
- **Midfielders** intercept threats and relay information, optimising interception, relay, and support objectives while bridging the defensive and offensive layers;
- **Forwards** project the mission outward, engaging threats under HO-CBF safety constraints.

This partition decomposes a high-dimensional multi-objective coordination problem into coupled lower-dimensional subproblems, one per role class. Enabling technologies include geometric rigid-body modelling on SE(3), CI-based distributed estimation that preserves consistency under unknown cross-correlations and intermittent connectivity, a decentralised HO-CBF quadratic-program safety filter, and a communication-aware autonomy mode for continued operation under degraded links.

1.4 Positioning against related role-heterogeneous work

Role and hierarchy in UAV swarms are not new in the abstract. Leader–follower architectures assign differentiated roles implicitly, and hierarchical heterogeneous schemes have been explored for formation control using distributed model predictive control with leader/coordinator/follower roles [42], for LLM-driven role-adaptive navigation [43], for role-decoupled reinforcement learning in aerial combat [44], and for cluster-head/member hierarchies in spectrum-monitoring swarms [45]. The present work differs in what the roles *are* and *how they are enforced*: SFSF roles are spatially layered (inner defensive shell, transitional midfield zone, outer engagement perimeter) with role-conditioned potential functions derived from screening geometry, integrated within a single unified architecture with CI estimation, HO-CBF safety filtering, dwell-time-constrained reassignment, and EW-aware degraded-mode operation. The contribution is a **structural integration and formalisation** of role-partitioned organisation for EW resilience, rather than the introduction of role differentiation itself.

1.5 Contributions and roadmap

The contributions are:

1. The RPSI paradigm: role-partitioned coordination formalised as an organisational principle, with a communication-cost and resource-allocation analysis (Section 3).
2. The SFSF framework: a three-layer role architecture with suitability-score assignment, quota constraints, and dwell-time-constrained transitions (Section 4).
3. A communication-resilient distributed estimation architecture based on CI fusion with a consistency guarantee (Section 6).
4. A decentralised HO-CBF safety filter integrated with role-conditioned potential control, with forward invariance of the safe set (Section 7).
5. A switched-system stability analysis establishing practical stability under average dwell-time role transitions (Section 8).
6. An extensive SITL evaluation with comparative baselines, five graded EW scenarios, scalability to $N = 100$, a 500-trial safety evaluation, and six ablations (Section 9), complemented by a reduced-order sensitivity study (Section 9.5).

Section 2 reviews related work; Section 10 discusses the results; Section 11 presents the limitations; Section 12 concludes. Proofs are in Appendix A.

2. Related Work

2.1 Consensus, leader–follower, and virtual structures

Consensus-based coordination [1–4] is distributed and scalable but requires persistent communication to preserve convergence, and typically treats agents as functionally equivalent. Leader–follower and virtual-structure methods [10,11] provide structure at the cost of leader vulnerability, since leader loss or link failure degrades the whole swarm, a critical weakness under EW. These observations motivate resilience mechanisms that do not depend on critical nodes.

2.2 Bio-inspired, optimisation-based, and learning-based coordination

Bio-inspired methods (Boids [6], PSO [8], ACO [9]) generate complex behaviour from simple rules but generally lack formal safety and convergence guarantees. Distributed optimisation and model predictive control [32,33] handle constraints explicitly but carry computational and communication burdens that grow with swarm size and depend on reliable state exchange. Multi-agent reinforcement learning [34,35] adapts behaviour through experience but typically lacks certified safety, and training demands substantial resources. Recent role-heterogeneous learning frameworks [43–45] demonstrate that role differentiation improves performance and reduces communication, corroborating the RPSI hypothesis from a learning-based direction; RPSI pursues the same structural hypothesis with analytically tractable control, estimation, and safety components.

2.3 Communication resilience and EW-aware coordination

Recent work addresses jamming detection and avoidance [38], topology and power optimisation under jamming [39], game-theoretic relay selection under communication and navigation jamming [40], and MARL-based anti-jamming channel and power management [41]. These studies optimise network-level decisions and do not reorganise mission-level responsibilities; relay selection and cluster-head policies provide functional differentiation but not the layered spatial–functional partition examined here. Communication impairment has also been studied in formation trajectory planning with explicit delay modelling [46]. The present work is complementary: robustness is achieved partly through estimation and control mechanisms and partly through the role decomposition itself.

2.4 Research gap

Existing work supports the following research-gap statement: existing swarm architectures either (i) assume functional homogeneity and therefore force every agent to reconcile competing objectives, (ii) differentiate roles without a spatially layered mission organisation integrated with consistent distributed estimation and certified safety filtering, or (iii) address communication degradation at the network layer without reorganising mission responsibilities. The literature reviewed here does not identify prior work that unifies role-partitioned spatial organisation, CI-based consistent estimation, HO-CBF safety filtering, and EW-aware degraded-mode autonomy in a single evaluated framework.

3. Role-Partitioned Swarm Intelligence: Formal Foundations

3.1 Definitions

Let the swarm consist of N agents with index set $\mathcal{N} = \{1, \dots, N\}$.

Definition 1 (Role partition). A role partition is a mapping $\Pi(t) : \mathcal{N} \to \mathcal{P}$, where $\mathcal{P} = \{D, M, F\}$ is the role set. At every time t, the swarm is partitioned into pairwise-disjoint subsets with $\mathcal{N}_D(t) \cup \mathcal{N}_M(t) \cup \mathcal{N}_F(t) = \mathcal{N}$ and $|\mathcal{N}_D| + |\mathcal{N}_M| + |\mathcal{N}_F| = N$. Agents in different subsets are governed by distinct objective functions and policies.

3.2 Objective decomposition

Let Γ denote the global mission objective. A homogeneous architecture minimises a weighted sum of all objective terms with a common policy for every agent. Under RPSI the global objective is decomposed into role-conditioned objectives $\Gamma = \sum_{r \in \mathcal{P}} J_r$, where J_D governs asset protection, J_M interception and support, and J_F offensive engagement. This transforms one high-dimensional problem into coupled lower-dimensional problems solved locally within each role class.

3.3 Communication-cost analysis

Counting dense exchanges within each role and between the Midfielder layer and the other layers gives the proxy

$$C_R = N_D^2 + N_M^2 + N_F^2 + N_M(N_D + N_F).$$

Proposition 1 (communication-cost bound). Under balanced roles, $N_D = N_M = N_F = N/3$, giving $C_R = 5N^2/9$,

compared with leading-order homogeneous cost $C_H \sim N^2$, a 44.4% constant-factor reduction in the dense-link model. If each agent instead has a maximum degree Δ independent of N, then $C_R \leq \Delta N = O(N)$, and $C_R/C_H \to 0$. The proof is given in Appendix A.

3.4 Resource-allocation efficiency

Let A be a fixed total action budget over a common horizon. Define objective interference as the weighted sum of positive-part alignments over simultaneously active objective-gradient pairs, $I_X = \sum_{q \in \mathcal{Q}_X} w_q [-\hat{g}_{q,1}^\top \hat{g}_{q,2}]_+$, and effective actions $E_X = A - \kappa I_X$, with $\kappa > 0$ and efficiency $\eta_X = E_X/A$.

Proposition 2 (resource-allocation efficiency). With a common action budget and horizon, if RPSI activates a subset of the objective pairs used by the homogeneous controller for every agent, then $\eta_R \geq \eta_H$. The proof is given in Appendix A.

3.5 Role transition dynamics

Role transitions are governed by a switching operator subject to a minimum dwell time $\tau_D > 0$, preventing excessive switching and supporting the stability analysis in Section 8.

4. The Soccer Field Swarm Formation Framework

4.1 Architecture

Consider a defended region containing protected assets $\Omega_{\text{assets}} = \{g_1, \dots, g_K\}$ and detected adversarial entities $\Omega_{\text{adv}}(t) = \{h_1, \dots, h_M\}$. SFSF organises agents into three concentric zones: an inner Defender layer (asset shielding, perimeter maintenance, last-line interception, gap minimisation), an intermediate Midfielder layer (threat interception, formation support, information propagation, dynamic reinforcement), and an outer Forward layer (threat neutralisation, early engagement, forward reconnaissance).

Figure 1 illustrates the spatial organisation of the three role layers.

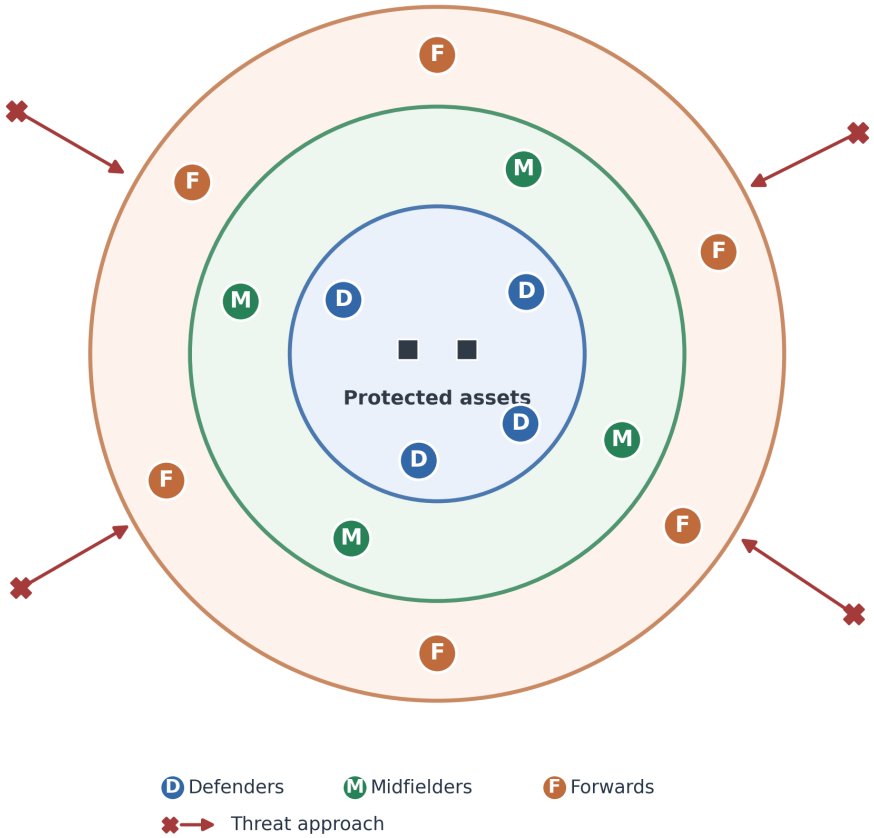


Figure 1. Conceptual SFSF spatial organisation. Defenders occupy the inner protection layer, Midfielders the intermediate interception and relay layer, and Forwards the outer engagement layer.

4.2 Role assignment

For each agent i and candidate role $r \in \{D, M, F\}$, the suitability score combines normalised energy, threat-related suitability, positional, and sensing features:

$$C_i^{(r)}(t) = w_E^{(r)} \tilde{E}_i(t) + w_T^{(r)} \tilde{T}_i^{(r)}(t) + w_P^{(r)} \tilde{P}_i^{(r)}(t) + w_S^{(r)} \tilde{S}_i^{(r)}(t).$$

The weights satisfy $w_k^{(r)} \geq 0$ and $\sum_k w_k^{(r)} = 1$. Features are mapped to $[0, 1]$ using predefined bounds, and the weights are held fixed during evaluation. Assignment uses binary variables $x_{ir} \in \{0, 1\}$, where $x_{ir} = 1$ indicates assignment of agent i to role r, with $\sum_r x_{ir} = 1$ and per-role quotas $N_r^{\min} = \lceil \rho_r^{\min} N \rceil \leq \sum_i x_{ir} \leq \lfloor \rho_r^{\max} N \rfloor = N_r^{\max}$, with a feasible quota configuration satisfying $\sum_r \rho_r^{\min} \leq 1 \leq \sum_r \rho_r^{\max}$. The feasible role assignment with the highest aggregate suitability is selected. A dwell-time lock holds the previous role assignment, $x_{i,r_i(t^-)}(t) = 1$, while $t - t_i^{\text{last}} < \tau_D$.

4.3 Role objectives

$$J_D = J_{\text{asset}} + J_{\text{coverage}} + J_{\text{intercept}},$$
$$J_M = J_{\text{support}} + J_{\text{relay}} + J_{\text{intercept}},$$
$$J_F = J_{\text{attack}} + J_{\text{recon}} + J_{\text{projection}}.$$

4.4 Communication architecture

Three levels operate hierarchically: intra-role, inter-role (Midfielder-mediated), and a degraded mode. When the local communication-health score falls below a threshold ($\lambda_i(t) < \lambda_{\text{crit}}$), the agent transitions to autonomous resilience operation: the role assignment is frozen and target tracking continues on local sensing alone, providing graceful degradation rather than mission failure.

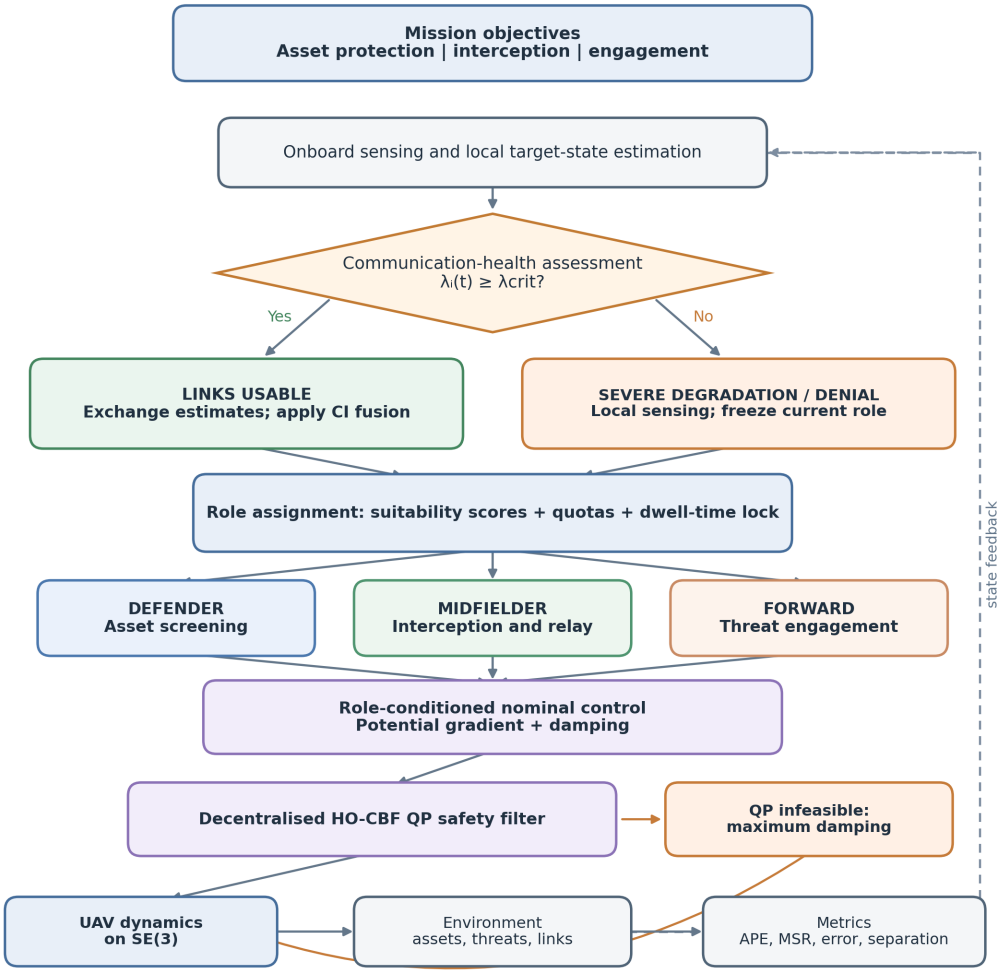


Figure 2. Overview of the RPSI-SFSF coordination and resilience architecture. Communication-health assessment selects CI-assisted fusion or local-only resilience operation; role-conditioned control is passed through an HO-CBF quadratic-program safety filter before actuation. The dashed arrow represents feedback to the next sensing and estimation cycle.

5. System Dynamics on SE(3) and Role-Conditioned Control

5.1 Rigid-body dynamics

To avoid Euler-angle singularities, each agent's state is $\xi_i = (p_i, v_i, R_i, \Omega_i) \in SE(3)$ with $p_i \in \mathbb{R}^3$, $v_i \in \mathbb{R}^3$, $R_i \in SO(3)$, $\Omega_i \in \mathbb{R}^3$. Translational dynamics:

$$m_i \dot{v}_i = m_i g e_3 - f_i R_i e_3 - D_\gamma v_i + u_i,$$

where m_i is mass, f_i collective thrust, g gravitational acceleration, D_γ translational drag, and u_i the swarm-level control input. Rotational dynamics:

$$\dot{R}_i = R_i \Omega_i^\times, \qquad J_i \dot{\Omega}_i = M_i - \Omega_i \times (J_i \Omega_i) - D_\Omega \Omega_i,$$

with J_i denoting inertia, M_i the moment input, D_Ω rotational damping, and $\Omega_i^\times$ the skew-symmetric map. SE(3) provides a globally valid representation; remaining limitations (modelling error and unmodelled aerodynamics) are discussed in Section 11.

5.2 Role-conditioned control laws

The swarm-level input decomposes as $u_i = u_i^{\text{role}} + u_i^{\text{safe}} + u_i^{\text{track}}$ with

$$u_i^{\text{role}} = -\nabla_{p_i} \Phi_i^r - D_c v_i, \qquad r \in \{D, M, F\},$$

Here, D_c is positive-definite damping and Φ_i^r is the role potential. Target assignment is the nearest currently estimated adversarial target, $j_i^r = \arg\min_j \|p_i - \hat{p}_j\|$, with $r \in \{M, F\}$ and ties broken by index. For Defenders, with assigned asset g_{k^*} and nearest-to-asset threat j_i^D, the screening point is

$$q_i^D = g_{k^*} + \rho_D(\hat{p}_{j_i^D} - g_{k^*}), \qquad 0 < \rho_D < 1,$$

lying on the asset–threat segment.

Defender potential.

$$\Phi_i^D = \frac{\alpha_1}{2}\|p_i - q_i^D\|^2 + \alpha_2 \sum_{\ell \in \mathcal{N}_D \setminus \{i\}} \frac{1}{\|p_i - p_\ell\|^{\nu_1}} + \alpha_3 \sum_k \frac{1}{\|p_i - g_k\|^{\nu_2}}.$$

Midfielder potential.

$$\Phi_i^M = \frac{\beta_1}{2}\|p_i - \hat{p}_{j_i^M}\|^2 + \frac{\beta_2}{2} \sum_{f \in \mathcal{F}_i(t)} \|p_i - p_f\|^2 + \beta_3 \sum_k \frac{1}{\|p_i - g_k\|^{\nu_3}}.$$

with $\mathcal{F}_i(t)$ denoting the set of Forwards supported by Midfielder i.

Forward potential.

$$\Phi_i^F = \frac{\mu_1}{2}\|p_i - \hat{p}_{j_i^F}\|^2 + \mu_2 \sum_{\ell \in \mathcal{N}_F \setminus \{i\}} \frac{1}{\|p_i - p_\ell\|^{\nu_4}} + \mu_3 \sum_k \frac{1}{\|p_i - g_k\|^{\nu_5}}.$$

The quadratic target-attraction term directs Forwards towards the currently estimated target. Gains $\alpha, \beta, \mu \geq 0$ are strictly positive on active terms; exponents $\nu > 0$ set repulsion steepness; all are fixed per experiment. The HO-CBF filter of Section 7 independently enforces inter-agent safety when it modifies the nominal control.

6. Communication-Resilient Distributed Estimation

6.1 Local estimation

Each target h has state $x_h = [p_h^\top, v_h^\top]^\top$ with dynamics $\dot{x}_h = Ax_h + w_h,\ w_h \sim \mathcal{N}(0, Q_h)$ [23,25]. Each UAV runs a prediction–correction filter from local observations z_k:

$$\begin{aligned} \hat{x}_{k|k-1} &= A\hat{x}_{k-1|k-1}, & P_{k|k-1} &= AP_{k-1|k-1}A^\top + Q, \\ K_k &= P_{k|k-1}H^\top(HP_{k|k-1}H^\top + R)^{-1}, & & \\ \hat{x}_{k|k} &= \hat{x}_{k|k-1} + K_k(z_k - H\hat{x}_{k|k-1}), & P_{k|k} &= (I - K_kH)P_{k|k-1}. \end{aligned}$$

6.2 Covariance Intersection fusion

To eliminate information loops and double counting in cyclic networks, estimates are fused in information form [18–20]:

$$Y_{\text{CI}} = \omega Y_i + (1-\omega)Y_j, \qquad y_{\text{CI}} = \omega y_i + (1-\omega)y_j,$$

with $Y_i = P_i^{-1}$ and $y_i = Y_i\hat{x}_i$; recovery via $P_{\text{CI}} = Y_{\text{CI}}^{-1}$ and $\hat{x}_{\text{CI}} = P_{\text{CI}}y_{\text{CI}}$.

Theorem 1 (CI consistency). If P_i and P_j are consistent covariances for two agents observing the same target, P_{CI} is conservative and consistent for all $\omega \in [0, 1]$ [18–20].

6.3 EW attenuation and mode toggling

Jamming is represented through packet reception probability $P_{\text{rx}} = (1 - J(t)) \exp(-\alpha d_{ij})$, with $J(t) \in [0, 1]$. The local communication-health score is $\lambda_i(t) = \sum_j w_{ij} P_{\text{rx}}$. If $\lambda_i(t) < \lambda_{\text{crit}}$, the agent enters autonomous resilience mode (Section 4.4), freezes its role assignment, and continues tracking with locally available estimates.

7. Safety-Critical Control via High-Order Control Barrier Functions

7.1 Safe set

Let $R_{\text{safe}} > 0$ be the minimum allowable inter-agent distance. The safe set is $\mathcal{C} = \{x : h_{ij}(x) \geq 0,\ \forall i \neq j\}$, with the relative-degree-two barrier

$$h_{ij}(x) = \|p_i - p_j\|^2 - R_{\text{safe}}^2.$$

Differentiating twice exposes the acceleration-level control dependence:

$$\begin{aligned}\dot{h}_{ij} &= 2(p_i - p_j)^\top (v_i - v_j),\\ \ddot{h}_{ij} &= 2\|v_i - v_j\|^2 + 2(p_i - p_j)^\top (\dot{v}_i - \dot{v}_j).\end{aligned}$$

7.2 Decentralised QP safety filter

Defining $\psi_{ij} = \dot{h}_{ij} + \alpha_1 h_{ij}$ and imposing $\dot{\psi}_{ij} + \alpha_2 \psi_{ij} \geq 0$ yields the per-agent quadratic program [27]

$$\begin{aligned}\min_{u_i} \quad & \frac{1}{2}\|u_i - u_i^{\text{nom}}\|^2\\ \text{s.t.} \quad & A_{\text{cbf}} u_i \leq b_{\text{cbf}}.\end{aligned}$$

If the QP is infeasible, a fallback applies maximum damping deceleration $u_i = -D_{\max} v_i$ to dissipate kinetic energy rapidly.

Theorem 2 (forward invariance). Consider the swarm dynamics with the HO-CBF constraints above. If the QP remains feasible at every update and all agents apply the resulting inputs, the safe set $\mathcal{C}$ is forward invariant [24,26].

8. Stability of the Switched Role-Partitioned Swarm

Closed-loop dynamics are modelled as a switched system $\dot{x} = f_{\sigma(t)}(x)$, where σ(t) tracks the active role configuration [30,31]. For a fixed configuration, the energy-based candidate

$$V(x) = \frac{1}{2}\sum_i v_i^\top v_i + \sum_i \Phi_{r_i}(p_i).$$

satisfies $\dot{V} = -\sum_i v_i^\top D_c v_i \leq 0$, so fixed role configurations are stable; under standard LaSalle arguments the swarm converges to stationary points of the role potentials [28,29]. Role switches introduce a bounded potential shift $|V^+ - V^-| \leq \Delta V$. Limiting the switching frequency by an average dwell-time condition $N_\sigma(T) \leq N_0 + T/\tau_D$ ensures that continuous energy dissipation dominates discrete switching-induced gains [31].

Theorem 3 (stability of RPSI). Consider the RPSI swarm with role-conditioned potentials, positive-definite damping, average dwell-time switching, and feasible HO-CBF constraints. Under the assumptions detailed in Appendix A, the closed-loop trajectories converge to a neighbourhood of the equilibrium set.

9. Experimental Evaluation

9.1 Platform and protocol

The primary evaluation uses a Software-in-the-Loop (SITL) platform coupling PX4 autopilot instances to a 6-DoF rigid-body engine at 240 Hz [36,37], with communication, sensing, and EW models. It comprises five graded communication-health scenarios (Table 1), comparative baselines under nominal conditions (Table 2), scalability tests to N = 100 (Table 3), EW degradation sweeps (Table 4), a 500-trial safety evaluation (Table 5), six ablations (Table 6), and onboard timing (Table 7). Section 9.5 presents a separate reduced-order packet-drop experiment.

Table 1. Communication-health degradation scenarios.

Scenario	Packet availability	Description
C1	100%	Nominal operation, complete link availability
C2	75%	Light interference and packet loss
C3	50%	Moderate adversarial jamming, fragmentation
C4	25%	Severe EW disruption
C5	0%	Complete communication denial

Metrics. Asset-Protection Effectiveness (APE) is the mean composite mission score formed from the fraction of threats neutralised and the fraction of assets kept uncompromised. Mission Success Rate (MSR) is the percentage of trials in which all threats are neutralised with no asset compromise. Estimation error is the mean target-position error of the fused estimate. The nominal inter-agent protection distance is $R_{\text{safe}} = 5.0\,\text{m}$. The reduced-order APE analogue uses equal weights for threat neutralisation and asset protection.

9.2 Comparative performance under nominal communication (N = 50)

Table 2. Mean ± std over trials.

Method	APE (%)	MSR (%)	Collisions	Est. error (m)
B1 Consensus control	71.4 ± 2.9	68.1 ± 3.4	11.2 ± 2.1	5.84 ± 0.71
B2 Leader–follower	74.9 ± 2.7	72.6 ± 3.1	9.5 ± 1.9	5.27 ± 0.66
B3 APF swarm	78.8 ± 2.4	76.1 ± 2.8	8.1 ± 1.6	4.92 ± 0.62
B4 APF + CBF	82.3 ± 2.1	80.7 ± 2.5	0.0 ± 0.0	4.66 ± 0.59
B5 Distributed MPC	86.1 ± 1.8	84.9 ± 2.1	0.0 ± 0.0	4.11 ± 0.53
RPSI-SFSF (proposed)	92.7 ± 1.4	91.2 ± 1.7	0.0 ± 0.0	3.24 ± 0.41

Under nominal communication, RPSI-SFSF has the highest APE (92.7 ± 1.4%) and mission success rate (MSR; 91.2 ± 1.7%) among the compared methods. Relative to distributed MPC, the differences are +6.6 percentage points in APE and +6.3 points in MSR. Fused target-position error is 3.24 ± 0.41 m for RPSI-SFSF and 4.11 ± 0.53 m for distributed MPC. RPSI-SFSF, APF + CBF, and distributed MPC each report zero collisions in this table, whereas the three methods without a certified safety filter report nonzero collision counts. These results show higher mean APE and MSR, and lower target-position error, for RPSI-SFSF than for the listed baselines under the tested nominal conditions.

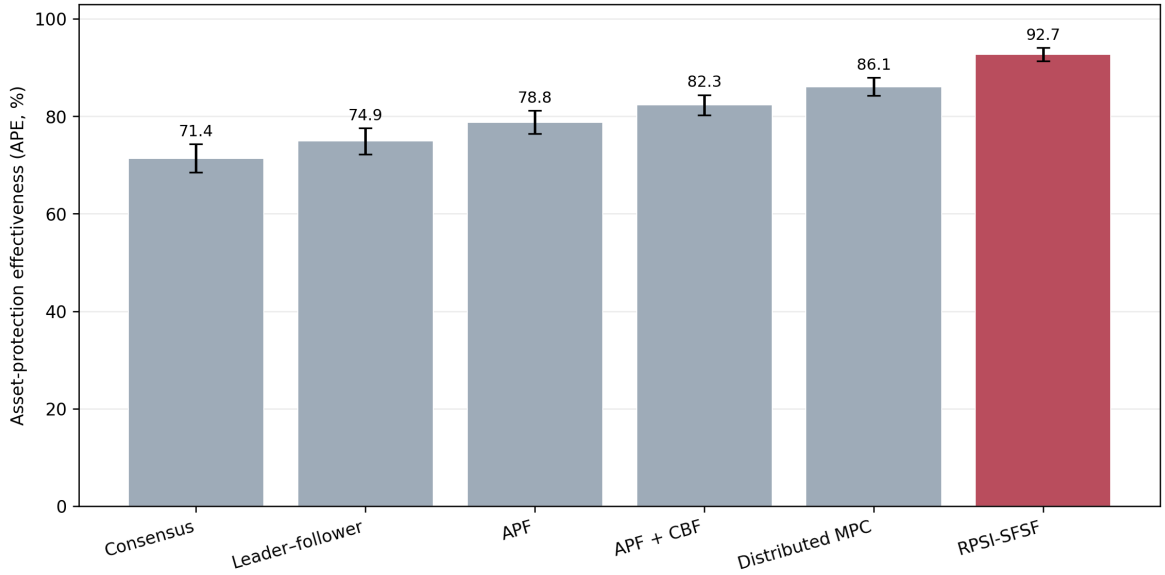


Figure 3. Nominal SITL APE comparison (Table 2). Bars show mean values and error bars show standard deviations.

9.3 Scalability and EW degradation

Table 3. Scalability (APE / MSR / comm. load / latency).

N	APE (%)	MSR (%)	Load (kbps/UAV)	Latency (ms)
10	95.4	94.8	26.1	14.2
25	94.1	93.0	32.7	17.6
50	92.7	91.2	39.3	22.8
75	91.8	89.6	44.5	28.4
100	90.9	88.7	49.8	33.1

Across N = 10 to N = 100, APE decreases from 95.4% to 90.9% (4.5 percentage points), while MSR decreases from 94.8% to 88.7% (6.1 points). Communication load per UAV rises from 26.1 to 49.8 kbps, and latency rises from 14.2 to 33.1 ms. Under the dense-link assumption, the analytical communication cost remains quadratic; lower asymptotic growth requires a bounded-degree topology.

Table 4. APE (%) under EW degradation.

Packets	Consensus	APF+CBF	Dist. MPC	RPSI-SFSF
100% (C1)	71.4	82.3	86.1	92.7
75% (C2)	67.9	80.5	84.3	91.1
50% (C3)	60.8	75.6	79.4	88.5
25% (C4)	49.7	66.2	71.8	84.2
0% (C5)	33.4	51.1	57.6	79.8

At C5, RPSI-SFSF achieves an APE of 79.8%, compared with 57.6% for distributed MPC, a difference of 22.2 percentage points (38.5% relative to the distributed-MPC value). Across C1 to C5, RPSI-SFSF decreases by 12.9 APE points (92.7% to 79.8%), whereas distributed MPC decreases by 28.5 points (86.1% to 57.6%). Figure 4 compares the trend across all five communication-health settings. RPSI-SFSF exhibits a smaller APE decline than distributed MPC across these simulated settings.

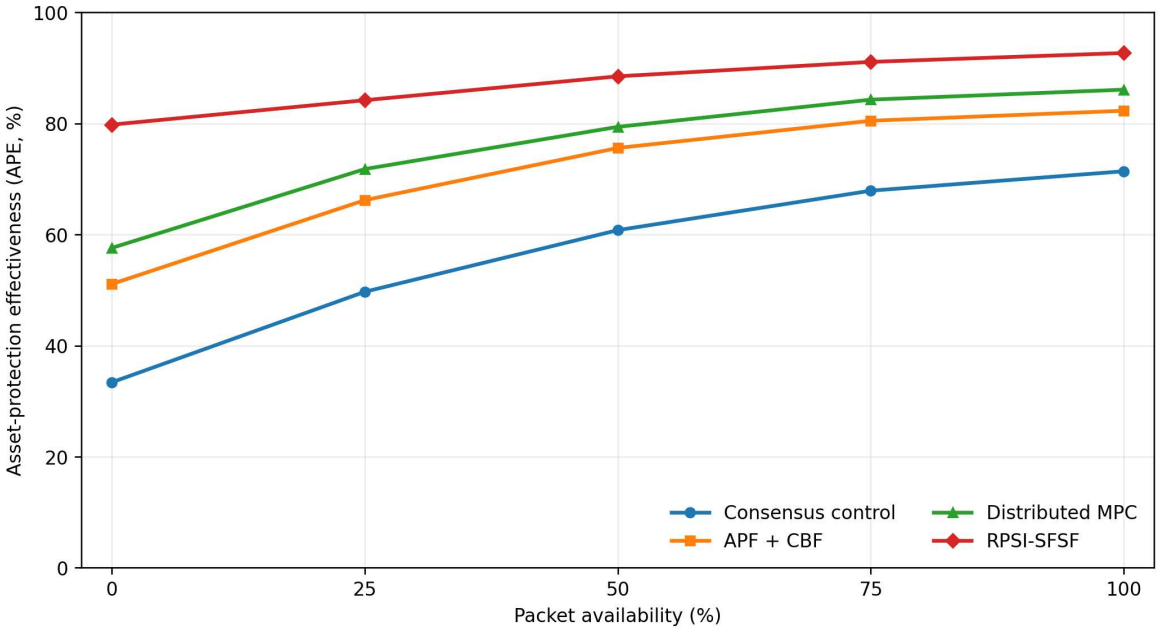


Figure 4. SITL APE across the five communication-health settings in Table 4.

9.4 Safety, ablations, and computation

Table 5. Safety evaluation (500 trials).

Metric	Value
Total inter-agent collisions	0
Near-miss events	3
Minimum recorded separation	5.37 m
Required protection distance R_safe	5.0 m
HO-CBF filter activation frequency	4.8%

Across the 500 SITL safety trials, the evaluation records zero inter-agent collisions, three near-miss events, a minimum separation of 5.37 m, and HO-CBF activation in 4.8% of control updates. The minimum observed separation is 0.37 m above the 5.0 m design criterion. The nominal role potentials maintained separation in most sampled updates, and the safety filter intervened when its constraints required correction.

Table 6. Ablations (nominal conditions).

Variant	APE (%)	MSR (%)
A1 Full framework	92.7	91.2
A2 Without role partitioning	83.6	81.9
A3 Without CI fusion	88.1	86.7
A4 Without HO-CBF filter	91.5	90.3
A5 Without comm-aware autonomy	85.8	83.4
A6 Without dynamic reassignment	87.6	85.2

Figure 5 compares each one-component ablation with the full framework. Removing role partitioning reduces APE by 9.1 percentage points, and removing communication-aware autonomy reduces it by 6.9 points. Removing dynamic reassignment and CI fusion reduces APE by 5.1 and 4.6 points, respectively. Removing the HO-CBF filter changes nominal APE by -1.2 points; APE captures mission effectiveness, while the filter enforces the separation constraints.

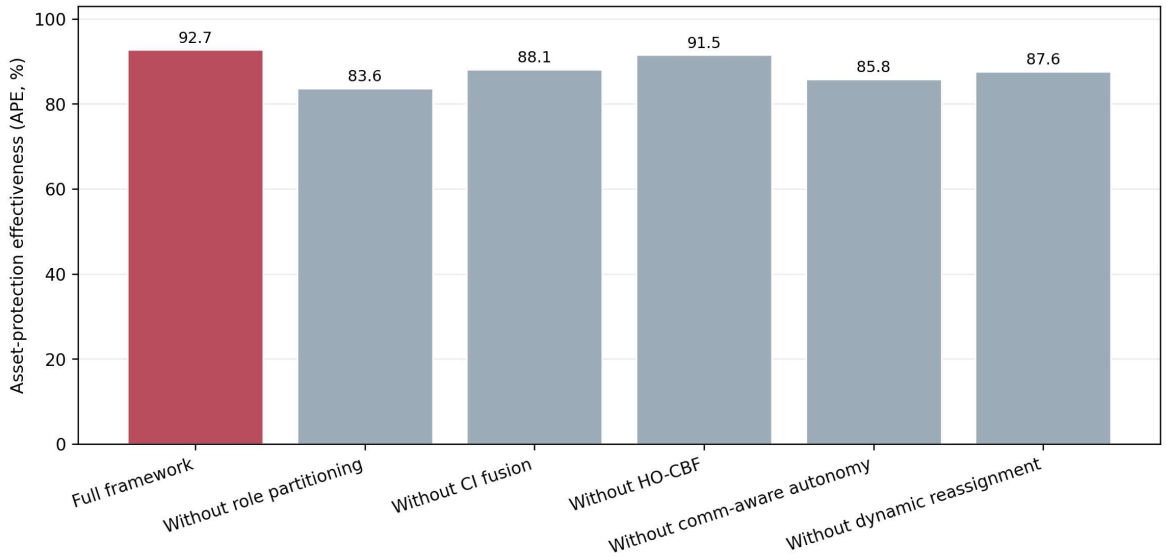


Figure 5. Nominal APE in the full framework and five component ablations (Table 6).

Table 7. Onboard computation.

Module	Mean time (ms)
Role assignment	1.4
CI fusion	2.1
HO-CBF QP	4.3
Communication loop	1.0
Total control cycle	8.8

The module means sum to 8.8 ms (1.4 + 2.1 + 4.3 + 1.0 ms), below the 10 ms period associated with a 100 Hz update rate.

9.5 Reduced-order packet-drop evaluation

A 2D point-mass simulation evaluates the sensitivity of role partitioning to target-state update packet loss and compares inter-agent separation under both controllers. The model uses $N = 12$ agents, four moving threats, two protected assets, and 50 s of simulated time at a 0.1 s step. For each controller and packet-delivery probability $p_{\text{packet}} \in \{1.0, 0.5, 0.0\}$, 20 seed-paired runs were completed, giving 120 controller runs. Paired controller runs use the same initial conditions, packet-delivery draws, observation-noise draws, and exogenous threat disturbances.

Each agent begins with a noisy target estimate. At every subsequent step, each agent-target update packet is delivered with independent Bernoulli probability p_{packet}. A received packet refreshes the estimate using the current target position plus Gaussian observation noise; a dropped packet leaves the last received estimate unchanged. Thus, 50% represents intermittent update delivery and 0% represents no subsequent target-state update packets during a run. This abstraction isolates the effects of stale target estimates; it does not model radio-frequency propagation, multi-hop routing, or specific jamming waveforms.

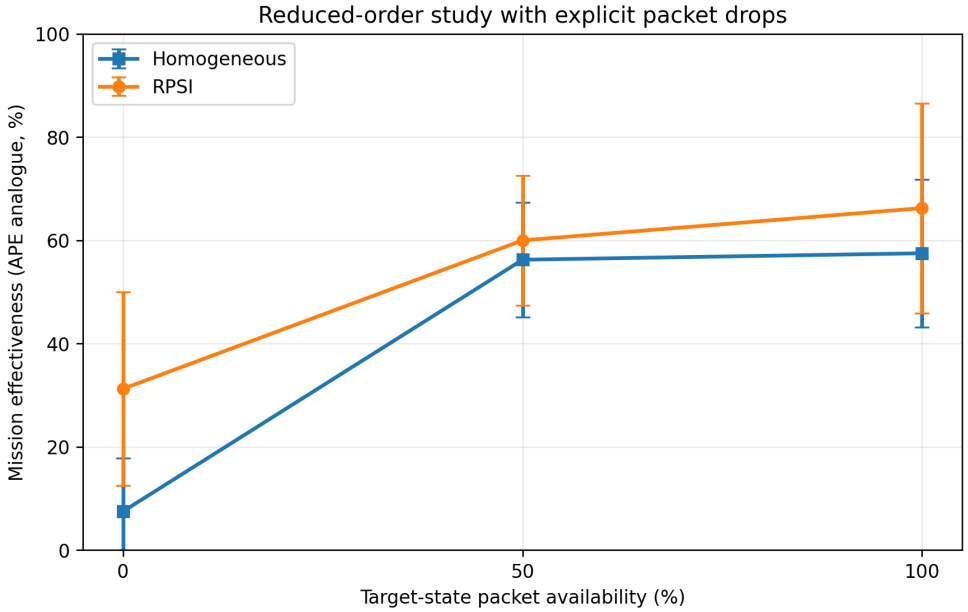


Figure 6. Reduced-order APE analogue as a function of target-state update packet-delivery probability. Error bars show one standard deviation across 20 runs per controller and setting.

Table 8. Reduced-order packet-drop results (20 runs per controller and setting).

Packet delivery	Controller	APE mean ± SD (%)	Minimum separation (m)	Runs below 5.0 m	Runs below 4.5 m
100%	RPSI	66.25 ± 20.32	5.00	0/20	0/20
100%	Homogeneous	57.50 ± 14.28	4.81	2/20	0/20
50%	RPSI	60.00 ± 12.57	4.92	4/20	0/20
50%	Homogeneous	56.25 ± 11.11	4.45	5/20	1/20
0%	RPSI	31.25 ± 18.80	5.00	0/20	0/20
0%	Homogeneous	7.50 ± 10.26	5.00	0/20	0/20

Mission-effectiveness results. RPSI exceeds the homogeneous controller's mean APE by 8.75 percentage points at 100% packet delivery, 3.75 points at 50%, and 23.75 points at 0%. The largest difference occurs when no target-state update packets arrive after initialization. The APE metric is quantised by the small numbers of threats and assets in this model, and across-run standard deviations range from 10.26 to 20.32 points. The observed difference therefore varies with packet availability and scenario outcome; the greatest gap in this experiment occurs under the zero-update condition.

Separation results. The minimum separation across all runs is 4.45 m. In the 100% packet-delivery setting, two homogeneous-controller runs fell below the 5.0 m design threshold, while no RPSI runs did. At 50% delivery, the threshold was crossed in four RPSI runs and five homogeneous-controller runs; one homogeneous run fell below 4.5 m. Neither controller crossed 5.0 m in the 0% setting. The reduced-order model uses velocity-level pairwise projection and positional correction for separation, while the primary SITL architecture uses the HO-CBF quadratic program described in Section 7.

The reduced-order experiment evaluates the effect of explicit target-state packet drops in a controlled paired-seed comparison. The findings show a larger role-partitioning advantage under the zero-update condition than under 50% delivery, while the separation results identify conditions that merit further safety tuning.

10. Discussion

Across the primary SITL scenarios, RPSI-SFSF achieves higher APE than the listed baselines, including a 22.2-percentage-point advantage over distributed MPC in the C5 setting. The ablations show the largest APE reductions when role partitioning and communication-aware autonomy are removed. In the reduced-order experiment with explicit target-state packet drops, RPSI exceeds the homogeneous controller by 8.75, 3.75, and 23.75 percentage points at 100%, 50%, and 0% packet delivery, respectively. Separation outcomes also depend on controller and delivery setting: some 100% and 50% runs fall below the 5.0 m design threshold. Together, the results indicate that role partitioning can improve mission effectiveness in the evaluated scenarios, with the magnitude depending on target-update availability and threat–asset geometry.

Comparison with the recent role-heterogeneous literature is informative. Learning-based role differentiation [43–45] achieves analogous communication reductions through training, whereas RPSI achieves them through structure, retaining interpretable potentials, consistency guarantees, and a certified safety filter at the cost of requiring hand-designed role objectives. Hierarchical DMPC with heterogeneous roles [42] shares the structural hypothesis but does not integrate CI-based EW-resilient estimation or HO-CBF safety in the same framework. The two approaches are complementary candidates for hybrid learning-augmented role adaptation (Section 12).

11. Limitations

1. **Simulation-only evidence.** All primary results are SITL. Wind, aerodynamic perturbation, sensor delay, non-Gaussian noise, and real RF propagation are not modelled; Section 9.5 adds a reduced-order study, not physical validation.
2. **Communication model.** The primary SITL EW model uses stochastic packet availability and exponential distance attenuation; the reduced-order experiment explicitly drops target-state update packets. Neither model reproduces RF propagation, correlated jamming bursts, spoofing, or adaptive adversaries.
3. **Conditional advantage.** The magnitude of the reduced-order mission-level advantage varies with packet delivery; the primary SITL scenario set does not systematically vary asset exposure or threat pressure.
4. **Quadratic communication growth.** Proposition 1 is a constant-factor result; sub-quadratic scaling requires bounded-degree topologies not evaluated here.
5. **QP feasibility.** Theorem 2 assumes feasibility; the fallback deceleration mitigates but does not certify behaviour in dense encounters.
6. **Parameter sensitivity.** Role weights, potential gains, and dwell time were fixed per experiment; a systematic sensitivity analysis is future work.
7. **Metric definitions and reproducibility.** APE and MSR are composite constructs; the reduced-order APE analogue is defined separately in Section 9.5. The primary SITL materials are available upon reasonable request.

12. Conclusions and Future Work

This paper presents Role-Partitioned Swarm Intelligence and its Soccer Field Swarm Formation implementation, combining role-conditioned objectives, Covariance Intersection-based estimation, an HO-CBF safety filter, and dwell-time-constrained role transitions. In the primary PX4 SITL evaluation, RPSI-SFSF reaches 92.7% APE under nominal communication compared with 86.1% for distributed MPC, and 79.8% compared with 57.6% in the C5 condition. In the 120-run reduced-order packet-drop experiment, mean APE is higher for RPSI by 8.75, 3.75, and 23.75 percentage points at 100%, 50%, and 0% target-state update delivery. The same experiment records a minimum separation of 4.45 m, with some 100% and 50% delivery runs falling below the 5.0 m design threshold. Future work will extend the evaluation to hardware-in-the-loop and flight experiments, examine richer adversarial communication models, and systematically vary asset exposure, threat pressure, and network topology.

Funding

This research received no external funding.

Data Availability Statement

The reduced-order summary dataset, the original PX4 SITL code, and the logs underlying Tables 2–7 will be available upon reasonable request.

Code Availability Statement

The reduced-order script, the figure-generation script, the architecture diagram source, and the original PX4 SITL implementation will be available upon reasonable request.

Author Contributions

M. H. S. Hassan: conceptualization, methodology, software, investigation, formal analysis, data curation, visualization, validation, writing—original draft, writing—review and editing, project administration. **B. J. Kumar:** software, investigation, validation, visualization. **N. S. M. Dudaboyina:** investigation, data curation, software. **G. R. S. Gunnu:** supervision, methodology review, writing—review and editing, project administration. **K. N. Rao:** supervision, academic guidance, writing—review and editing, project administration.

Generative AI Statement

Generative AI tools were used during manuscript preparation to refine language and improve readability and to assist in generating the manuscript's figures. They were not used to originate the research concept. The authors reviewed the final text and figures and take responsibility for the content of the manuscript.

Appendix A. Proofs

Proof of Proposition 1 (communication-cost bound). With balanced roles, $N_D = N_M = N_F = N/3$ and $C_R = 3(N/3)^2 + (N/3)(2N/3) = 5N^2/9$. The fully connected homogeneous baseline has $N(N-1)$ links, so its leading-order communication cost is $C_H \sim N^2$. Thus, $\lim_{N\to\infty} C_R/C_H = 5/9 < 1$ and $(C_H - C_R)/C_H \to 4/9$. If every intra-role and inter-role link set has maximum degree Δ independent of N, then each agent sends at most Δ messages per cycle, yielding $C_R \le \Delta N = O(N)$ and $C_R/C_H \to 0$. □

Proof of Proposition 2 (resource-allocation efficiency). Each pairwise interference term is nonnegative because the weights and positive-part operator are nonnegative. The RPSI interference sum covers a subset of the homogeneous pairs, so $I_R \le I_H$. With common action budget A, $\eta_R - \eta_H = \kappa(I_H - I_R)/A \ge 0$. □

Proof of Theorem 3. Between switches, $\dot{V} \le 0$, and LaSalle's invariance principle gives convergence to the largest invariant subset of $\{v_i = 0\}$ within the active role configuration, i.e., critical points of $\sum_i \Phi_{r_i}(p_i)$. At a switch, V increases by at most ΔV; under the average dwell-time bound $N_\sigma(T) \le N_0 + T/\tau_D$, the total injected energy over $[0, T]$ is bounded by $\Delta V(N_0 + T/\tau_D)$. Choosing τ_D sufficiently large for dissipation between switches to dominate these injections yields bounded energy and convergence to a neighbourhood of the equilibrium set, whose size shrinks with ΔV and $1/\tau_D$. Feasible HO-CBF constraints preserve the safe set $\mathcal{C}$ by Theorem 2. □

References

[1] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: Algorithms and theory," *IEEE Transactions on Automatic Control*, vol. 51, no. 3, pp. 401–420, 2006.

[2] R. Olfati-Saber, J. A. Fax, and R. M. Murray, "Consensus and cooperation in networked multi-agent systems," *Proceedings of the IEEE*, vol. 95, no. 1, pp. 215–233, 2007.

[3] W. Ren and R. W. Beard, "Consensus seeking in multi-agent systems under dynamically changing interaction topologies," *IEEE Transactions on Automatic Control*, vol. 50, no. 5, pp. 655–661, 2005.

[4] W. Ren and R. W. Beard, *Distributed Consensus in Multi-Vehicle Cooperative Control*. London, U.K.: Springer, 2008.

[5] M. Mesbahi and M. Egerstedt, *Graph-Theoretic Methods in Multiagent Networks*. Princeton, NJ, USA: Princeton University Press, 2010.

[6] C. W. Reynolds, "Flocks, herds and schools: A distributed behavioral model," *ACM SIGGRAPH Computer Graphics*, vol. 21, no. 4, pp. 25–34, 1987.

[7] T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, and O. Shochet, "Novel type of phase transition in a system of self-driven particles," *Physical Review Letters*, vol. 75, no. 6, pp. 1226–1229, 1995.

[8] J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of the IEEE International Conference on Neural Networks*, 1995, pp. 1942–1948.

[9] M. Dorigo and L. M. Gambardella, "Ant colony system: A cooperative learning approach to the traveling salesman problem," *IEEE Transactions on Evolutionary Computation*, vol. 1, no. 1, pp. 53–66, 1997.

[10] R. W. Beard and T. W. McLain, *Small Unmanned Aircraft: Theory and Practice*. Princeton, NJ, USA: Princeton University Press, 2012.

[11] N. Michael, J. Fink, and V. Kumar, "Cooperative manipulation and transportation with aerial robots," *Autonomous Robots*, vol. 30, no. 1, pp. 73–86, 2011.

[12] S.-J. Chung, A. A. Paranjape, P. Dames, S. Shen, and V. Kumar, "A survey on aerial swarm robotics," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 837–855, 2018.

[13] L. Bayındır, "A review of swarm robotics tasks," *Neurocomputing*, vol. 172, pp. 292–321, 2016.

[14] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, "Swarm robotics: A review from the swarm engineering perspective," *Swarm Intelligence*, vol. 7, no. 1, pp. 1–41, 2013.

[15] T. Lee, M. Leok, and N. H. McClamroch, "Geometric tracking control of a quadrotor UAV on SE(3)," in *Proceedings of the IEEE Conference on Decision and Control*, 2010, pp. 5420–5425.

[16] T. Lee, M. Leok, and N. H. McClamroch, "Geometric control of multiple quadrotor UAVs transporting a cable-suspended rigid body," *IEEE Transactions on Control Systems Technology*, vol. 26, no. 1, pp. 255–264, 2018.

[17] F. Bullo and A. D. Lewis, *Geometric Control of Mechanical Systems*. New York, NY, USA: Springer, 2005.

[18] S. Julier and J. Uhlmann, "A non-divergent estimation algorithm in the presence of unknown correlations," in *Proceedings of the American Control Conference*, 1997, pp. 2369–2373.

[19] S. Julier and J. Uhlmann, "General decentralized data fusion with covariance intersection," in *Handbook of Multisensor Data Fusion*. Boca Raton, FL, USA: CRC Press, 2001.

[20] M. B. Hurley, "An information theoretic justification for covariance intersection and its generalization," in *Proceedings of the International Conference on Information Fusion*, 2002.

[21] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation*. New York, NY, USA: Wiley, 2001.

[22] D. Simon, *Optimal State Estimation*. Hoboken, NJ, USA: Wiley, 2006.

[23] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Journal of Basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.

[24] A. D. Ames, X. Xu, J. W. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.

[25] A. D. Ames, S. Coogan, M. Egerstedt, K. Notomista, K. Sreenath, and P. Tabuada, "Control barrier functions: Theory and applications," in *Proceedings of the European Control Conference*, 2019.

[26] W. Xiao and C. Belta, "High-order control barrier functions," *IEEE Transactions on Automatic Control*, vol. 67, no.

7, pp. 3655–3662, 2022.

[27] A. Agrawal and K. Sreenath, "Discrete control barrier functions for safety-critical control," *IEEE Control Systems Letters*, vol. 1, no. 1, pp. 66–71, 2017.

[28] H. K. Khalil, *Nonlinear Systems*, 3rd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2002.

[29] J. P. LaSalle, *The Stability of Dynamical Systems*. Philadelphia, PA, USA: SIAM, 1976.

[30] D. Liberzon, *Switching in Systems and Control*. Boston, MA, USA: Birkhäuser, 2003.

[31] J. P. Hespanha and A. S. Morse, "Stability of switched systems with average dwell-time," in *Proceedings of the IEEE Conference on Decision and Control*, 1999, pp. 2655–2660.

[32] A. Richards and J. P. How, "Robust distributed model predictive control," *International Journal of Control*, vol. 80, no. 9, pp. 1517–1531, 2007.

[33] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.

[34] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: MIT Press, 2018.

[35] L. Busoniu, R. Babuška, and B. De Schutter, "A comprehensive survey of multi-agent reinforcement learning," *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, vol. 38, no. 2, pp. 156–172, 2008.

[36] V. Kumar and N. Michael, "Opportunities and challenges with autonomous micro aerial vehicles," *The International Journal of Robotics Research*, vol. 31, no. 11, pp. 1279–1291, 2012.

[37] S.-J. Chung, M. Dorothy, A. Paranjape, P. Dames, S. Shen, and V. Kumar, "Toward autonomous swarm robotics," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 3, pp. 409–435, 2020.

[38] S. Xing, S. Peccoud, S. Li, S. Li, and T. Yang, "Robust communication-aware jamming detection and avoidance for UAS swarms," *IEEE Access*, vol. 13, pp. 200431–200445, 2025, doi: 10.1109/ACCESS.2025.3630211.

[39] H. Li, H. Wang, X. Zheng, J. Gu, X. Guan, and H. Liu, "Topology optimization for UAV swarm communication with jamming," *IEEE Communications Letters*, vol. 29, no. 5, pp. 983–987, 2025, doi: 10.1109/LCOMM.2025.3549253.

[40] X. Ma, M. Yu, Z. Luo, and M. Gao, "Game-theoretic optimization of UAV swarm and relay communication for interference resilience," *IEEE Transactions on Vehicular Technology*, vol. 75, no. 3, pp. 4581–4594, 2026, doi: 10.1109/TVT.2025.3614715.

[41] Y. Li, T. Zhao, Z. Wu, Y. Lin, and Y. Zhang, "Energy-harvesting-assisted UAV swarm anti-jamming communication based on multi-agent reinforcement learning," *Drones*, vol. 10, no. 4, Art. no. 294, 2026, doi: 10.3390/drones10040294.

[42] J. Zhao, J. Sun, Z. Cai, Y. Wang, and K. Wu, "Distributed coordinated control scheme of UAV swarm based on heterogeneous roles," *Chinese Journal of Aeronautics*, vol. 35, no. 1, pp. 81–97, 2022, doi: 10.1016/j.cja.2021.01.014.

[43] Z. Wang, R. Li, S. Li, Y. Xiang, H. Wang, Z. Zhao, and H. Zhang, "RALLY: Role-adaptive LLM-driven yoked navigation for agentic UAV swarms," arXiv:2507.01378 (https://arxiv.org/abs/2507.01378), 2025.

[44] L. Wang, Y. Liu, L. Guo, R. Zhang, Y. Cui, Y. Cui, and C. Xu, "A cognitive synergetic hierarchical framework for UAV swarm combat via speculative inference and role-decoupled reinforcement learning," *Frontiers in Neurorobotics*, vol. 20, Art. no. 1799054, 2026, doi: 10.3389/fnbot.2026.1799054.

[45] Y. Sun, X. Zhang, M. Wang, Y. Yang, T. Xia, X. Zhu, and T. Cui, "Hierarchical role-based multi-agent reinforcement learning for UHF radiation source localization with heterogeneous UAV swarms," *Drones*, vol. 10, no. 1, Art. no. 54, 2026, doi: 10.3390/drones10010054.

[46] J. Yu, H. Fan, L. Wang, J. Gao, and Y. Yi, "DI-MATD3-based multi-UAV formation trajectory planning under asynchronous communication delays," *Drones*, vol. 10, no. 10, Art. no. 762, 2026. Available: https://www.mdpi.com/2504-446X/10/10/762.