

ORIGINAL RESEARCH

Open access

A Communication-State Adaptive Coordination Framework for UAV Swarm Tracking in Congested Networks

Hans Müller^{1*}, Anna Schmidt¹, Thomas Weber²

Abstract

Swarm-based unmanned aerial vehicle (UAV) systems are increasingly recognized for their capabilities in extensive IoT and precision agriculture operations, notably in collaborative surveillance and the tracking of multiple moving objects across outdoor terrains. Nevertheless, prevailing coordination mechanisms tend to assume flawless communication channels, disregarding practical network disturbances including congestion, dropped packets, and transmission delays. Such disruptions hinder the prompt relay of data among UAVs and the terrestrial control station, causing control instructions to arrive late or not at all. The outcome is a breakdown in coordination fidelity and a steep decline in tracking effectiveness when these systems are fielded. To remedy this shortcoming, we put forward CSOOC (Communication-State Driven Online–Offline Coordination with Congestion Control), a hybrid governance scheme that fuses centralized learning-powered decision processes with decentralized rule-governed behaviors, enabling UAVs to tailor their actions to instantaneous network status. The CSOOC framework is built around three core elements: (1) an online component providing centralized orchestration while communication remains dependable, (2) an offline, profit-guided movement policy relying on localized Gaussian maps to sustain autonomous target pursuit when connectivity drops, and (3) congestion mitigation rooted in STAR (Stratified Transmission and RTS/CTS), which couples staggered transmission timing with RTS/CTS handshake protocols to bolster uplink robustness. A cohesive co-simulation infrastructure is constructed, linking network-layer communication with swarm-level control and coordination dynamics. Trials reveal that CSOOC delivers a mean observation ratio of 39.7%, outperforming reference algorithms by margins of 4.4%–11.13%, while also boosting network dependability through markedly improved packet delivery success under saturated channel conditions. These outcomes confirm that CSOOC successfully spans the disconnect between algorithm performance within simulators and the operational realities of UAV collectives in communication-challenged settings.

Keywords UAVs, Multi-target tracking, Communication-aware, Congestion control

*Correspondence:

Hans Müller

hans.mueller@uni-bonn.de

¹ Department of UAV Swarm Coordination and Communication, Faculty of Engineering, University of Bonn, Bonn, Germany

² Department of Adaptive Network Protocols for Agricultural Robotics, Faculty of Engineering, LMU Munich, Munich, Germany

Introduction

Over the past several years, UAV collectives have exhibited notable promise within expansive IoT and intelligent farming domains, spanning activities such as vegetation surveillance, targeted agrochemical application, and ambient condition monitoring [1–3]. An essential function powering these activities is the joint observation of multiple dynamic or mobile entities across complicated landscapes—a challenge formally designated as the Cooperative Multi-Robot Observation of Multiple Moving Targets (CMOMMT) problem [4]. The principal aim in CMOMMT revolves around elevating the mean

observation percentage over the full set of targets through ongoing trade-offs between target pursuit and exploratory sweeping, working within each UAV's bounded sensing footprint and movement capabilities, as portrayed in **Figure 1**.




Figure 1. Visualization of a UAV collective executing cooperative target tracking within an IoT monitoring environment. Transmission bottlenecks between numerous UAVs and the GBS can trigger packet loss and delayed control loops.

A diverse collection of coordination approaches has been examined to confront the CMOMMT challenge, extending from early heuristic constructs to current learning-centric paradigms [4-8]. Pioneering techniques, exemplified by force-vector and frontier-guided schemes, employed localized rulesets for distributed teamwork. Later waves of optimization-driven strategies (including profit-maximization formulations, Particle Swarm Optimization, and Ant Colony Optimization) raised the efficiency of task distribution but demanded manually engineered cost functions. Reinforcement learning (RL) frameworks further boosted adaptability yet frequently grappled with prohibitive sample demands and stability concerns when deployed at scale.

Within the category of learning-centric approaches, our prior research delivered the Supervised and Transfer Learning Framework (STLF) [9], a dual-phase architecture wherein a Teacher Model is trained on privileged global state information and subsequently distills its knowledge into a Student Model that operates under incomplete situational awareness. STLF achieved considerable success in securing elevated observation percentages and surpassing multiple comparator approaches. That said, the baseline STLF presupposes an idealized setting of flawless, synchronous information exchange between UAVs and the Ground Base Station (GBS). Realistically, UAV swarms function within

wireless spectra featuring restricted capacity and contended uplink pathways. Extending the foundational STLF design, the present study elevates its viability for authentic networking scenarios through the incorporation of mechanisms that counteract communication breakdowns such as channel saturation, signal collisions, and data loss—phenomena routinely encountered in actual UAV swarm operations.

When numerous UAVs attempt concurrent transmission of state updates or receipt of directives from the GBS, channel congestion may materialize, generating packet erasures, response lag, or outright connectivity disruption [10, 11]. Given that centralized controllers like STLF presuppose prompt and trustworthy communication, such network deterioration directly undermines coordination quality and aggregate tracking outcomes. Investigations demonstrate that while Ultra-Wideband technology can furnish precise inter-UAV distance measurements, even broadband platforms face limitations from finite spectral resources and vulnerability to electromagnetic interference within densely packed swarms [12, 13]. These insights underscore that communication fidelity represents a foundational constraint, underscoring the necessity for coordination methods that are intrinsically sensitive to networking conditions. The situation demands a communication-cognizant swarm management framework that dynamically modulates UAV conduct in accordance with real-time channel status.

To rise to this challenge, we introduce CSOOC, a hybridized control architecture that blends learned and rule-driven behavioral policies while explicitly contending with communication variability. CSOOC incorporates three tightly interwoven sub-systems. In the first, during online operation, UAVs enact globally synchronized control directives issued by the STLF engine, capitalizing on centralized reasoning for productive target coverage. In the second, during offline intervals, UAVs revert to a utility-guided movement regime constructed from locally maintained Gaussian potential fields, empowering autonomous target shadowing and exploration amid communication blackouts. In the third, to facilitate fluid and dependable transitions between the two operational modes, a STAR congestion regulation unit is embedded within the system. Through the fusion of temporally dispersed transmission scheduling with RTS/CTS (Request-To-Send/Clear-To-Send) [14] channel pre-allocation, STAR dampens network congestion and fortifies uplink consistency, curtailing spurious mode oscillations and safeguarding steady control output.

The principal scholarly contributions delivered by this investigation are cataloged as follows:

- A communication-cognizant UAV collective coordination paradigm, designated CSOOC, is unveiled, supporting adaptive online–offline behavioral reconfiguration amidst authentic and fluctuating networking regimes.
- A utility-driven offline mobility protocol leveraging Gaussian spatial representations is formulated to preserve geographically informed target tracking throughout communication gaps.
- A bifurcated congestion suppression scheme, termed STAR, is engineered by unifying temporally layered transmission with RTS/CTS dialog to amplify network resilience and curtail packet contention.
- An integrated co-simulation environment is forged, merging networking and swarm control domains, thereby permitting methodical appraisal of coordination efficacy under genuine communication impediments.

These contributions additionally supply a fresh vantage point on coupling network emulation with swarm coordination dynamics, opening avenues for examining UAV conduct under constrained communications. STAR attenuates the consequences of channel saturation, while the online–offline switching construct empowers the collective to recalibrate its cooperative posture in response to present communication circumstances.

The balance of this manuscript proceeds as follows. The subsequent section surveys prior scholarship on CMOMMT coordination and communication-sensitive swarm governance. The following section formalizes the problem definition and system blueprint. The next section details the proposed CSOOC methodology. The subsequent section documents exhaustive empirical evaluations across diverse networking and swarming parametrizations. The final section synthesizes conclusions and charts prospective research trajectories.

Related Work

This segment surveys prior scholarship through four lenses, advancing from UAV swarm orchestration in CMOMMT settings toward communication bottlenecks, networked simulation environments, and congestion regulation techniques.

UAV swarm coordination for CMOMMT tasks

The CMOMMT challenge [4] strives to extend uninterrupted target visibility by means of joint multi-agent maneuvers. In the domain of UAV collectives, pioneering solutions like B-CMOMMT [15] and P-CMOMMT [16] resorted to rule-of-thumb and utility-weighted assignment schemes, striking a balance between pursuit and reconnaissance yet regularly delivering globally subpar results. Ensuing research brought in reconfigurable formation governance [17] and trajectory anticipation [5] to raise spatial coverage effectiveness.

Fresh scholarship has turned to Deep Reinforcement Learning (DRL) [18] to pick up flexible behavioral rules inside intricate surroundings, although such works routinely assume unbroken connectivity. Our antecedent STLF [9] architecture reached elevated observation percentages under flawless linkage but did not confront intermittent channels. This deficiency propels our embedding of communication-adaptive mode toggling for durable UAV swarm orchestration spanning both linked and disconnected regimes.

Communication-constrained multi-agent systems

Connectivity restrictions exert a decisive influence on teamwork within multi-agent constructs. Earlier formulations habitually took for granted pristine linkage [19], setting aside throughput caps, propagation lag, and channel volatility. Present-day inquiries fold communication cognizance into planning routines, such as reducing map resolution to enable lightweight live route computation [20]. For submerged AUV undertakings, vision-anchored deep reinforcement learning allows concurrent refinement of movement and signaling inside taxing acoustic conduits [19]. Throughout sprawling IoT and sensing fabrics, power-mindful and topology-shifting methodologies bolster connectedness under tight resource envelopes [21, 22]. These bodies of work make clear that fusing trajectory design with faithful communication depictions is indispensable for sturdy operation inside throughput- and delay-impoverted contexts.

Network simulation and communication

Dependable benchmarking of UAV transmission tactics calls for lifelike network emulation toolkits competent at reproducing channel fluctuations, nodal displacement, and protocol interplay. The NS-3 [23] simulation suite enjoys extensive uptake among wireless networking researchers thanks to its pluggable architecture and customizability. The emergence of LTE and 5G add-ons for NS-3 has empowered scholars to probe cellular-backed UAV links beneath faithful PHY/MAC-layer assumptions, absorbing mobility and collision effects [24, 25]. Extending beyond protocol-bounded modules, all-purpose emulation landscapes examined by Jiang *et al.* [26] provide direction on picking fitting test rigs for UAV contexts. High-fidelity emulation backends additionally accommodate pairing with movement simulators, permitting clock-synced scrutiny of both channel performance and UAV assignment metrics simultaneously [27]. These apparatuses aid the dissection of network operation across heterogeneous settings, from tightly packed metropolitan cores to uncluttered rural stretches, steering protocol tuning for UAV swarms. Contemporary publications accentuate the imperative of weaving in faithful propagation profiles, saturation phenomena, and displacement tracks to narrow the gulf separating simulation outcomes from on-site deployments [10].

Congestion control in UAV-ground communication

Agile congestion abatement stands as a cornerstone for UAV–earth linkages, given the fluid topology, pinched capacity, and brisk locomotion native to UAV meshes. Cross-stack blueprints, typified by the Mobility and Congestion-Aware Routing Protocol (MCARP), blend routing with transport-layer logic to responsively modulate traffic streams informed by movement forecasts and pathway fidelity, thereby cutting packet wastage and transit lag in shifting UAV nets [28]. Under TCP-grounded schools of thought, deep reinforcement learning has been co-opted to tune congestion window scaling

inside UAV-assisted radio infrastructures, eclipsing standard TCP lineages when faced with jittery link qualities [29]. Power-cautious pathways additionally unite congestion mitigation with measurement amalgamation schemes to reconcile network carriage against UAV energy depletion, elongating sortie duration across multi-UAV IoT ecosystems [30]. Chance-driven dispatch blueprints, akin to those put forward in load-sensitive multi-hop relaying, capitalize on buffer backlogs and channel condition readings to forestall chokepoints and lift delivery fractions beneath surging loads [31]. Moreover, congestion-perceptive air–ground composite constructions shift resource allotments fluidly between UAV relay stations and rooted nodes to dampen cross-talk and hold QoS across mixed topologies [32]. Aggregately, these inquiries bring into relief that meshing congestion mindfulness with motion extrapolation, cross-stack fine-tuning, and energetic frugality stands essential to upholding Hardy UAV–terrestrial ties beneath elevated traffic intensity.

Task Formulation

Mission overview and notation

Within this inquiry, we zero in on the CMOMMT formulation while nestling true-to-life network communication representations. By threading a network emulator (NS-3 [23]) through the operational pipeline, the choice loop accommodates transport holdups, datagram shedding, and capacity capping, thereby echoing more field-plausible UAV flight envelopes.

The diagram of the CMOMMT construct is rendered in Figure 2, with the fundamental constituents catalogued in what follows.




Figure 2. CMOMMT formulation [9].

Search Region. The scouting expanse is cast as a planar square lattice, sliced into equal-sized units delineated by gray stippled borders.

Time Step. The chronology is chopped into countable spans tagged as $t=1,2,\dots,T$, where T captures the aggregate endurance of the assignment.

Targets. A pool of M interchangeable subjects populates the arena. The whereabouts of subject j at interval t is written as (t) , pinpointing the grid compartment it inhabits. At every temporal tick, subjects advance into one of four side-adjacent compartments (upward, downward, leftward, or rightward), abiding by a circumscribed movement regime.

UAVs. A company of N interchangeable drones is scattered over the scouting expanse to keep watch on the M subjects ($N < M$). The drones command a superior ceiling speed relative to the subjects; bereft of this edge, a subject could handily shake off tailing by running at full tilt. In addition, drones hold a constant elevation across the whole sortie. Each drone is provisioned with:

1. GPS for position fixing.
2. Airframe-mounted visual sensors to register ground subjects. Drone i scans compartments falling within a stipulated Euclidean span, shaping its viewing envelope (t) . If subject j sits inside (t) , i.e., $CoTj(t) \in FoVi(t)$, it counts as spotted. Zones lying outside (t) persist uncharted to drone i , materializing fragmentary situation appreciation.
3. Radio gear to forward sightings to the GBS over a sculpted network (e.g., LTE/5G) inside NS-3, where signaling undergoes lag, packet dropouts, and saturation repercussions.

CMOMMT workflow

Figure 3 pictures the mission cadence. At each time interval, drones commence with localized sensing. The harvested readings are then shipped across the emulated network to the GBS. The GBS merges feeds from the full drone complement and derives harmonized navigation edicts, which are routed back to the drones through the same network. This loop, laid bare to communication lags and potential linkage snap-offs, reiterates until $t = T$.




Figure 3. CMOMMT Workflow with integrated network simulation.

Objective function

The goal is to maximize the mean detection fraction over the full set of subjects across T , signified as $\Theta(T)$:

$$\Theta(T) = \frac{1}{M * T} \sum_{t=1}^T \sum_{j=1}^M O_j(t) \quad (1)$$

Within which (t) stands as a two-state indicator encoding whether subject j falls under observation at instant t :

$$O_j(t) = \begin{cases} 1, & \text{if } \exists UAV_i : CoT_j(t) \in FoV_i(t) \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

Inside the platonic scenario where every subject gets detected at each time tick, $\Theta(T)=1$. In the field, however, given drone agility confines, tight FoV spans, and the fallout from network latency and packet shedding, hitting $O(t)=1$ for every j, t proves largely out of reach. The mission pursuit is to craft an orchestration rule that pushes $\Theta(T)$ upward while respecting both locomotor and connective shackles.

Target mobility model

To produce ecologically plausible movement sequences, we implement a Lévy flight-governed locomotion framework that reproduces the patchy, non-linear displacement signatures characteristic of animal foraging in nature [9]. Where classical uncorrelated random strolls tether agents within tight spatial confines, Lévy flights synthesize a spectrum of predominantly tiny hops punctuated by sporadic sweeping relocations, shaped by a distribution whose tails decay slowly according to a power law. This biologically grounded property elevates the problem's intricacy, demanding that the drone collective accommodate richly varied and unpredictable motion regimes. With each iteration i , the agent's spatial coordinates evolve according to

$$\mathbf{r}_i = \mathbf{r}_{i-1} + s_i \hat{\mathbf{d}}_i, \quad (3)$$

With s_i representing the displacement magnitude obtained through Lévy-distributed sampling and $\hat{\mathbf{d}}_i$ constituting a normalized directional vector. To honor velocity limitations imposed within the simulation environment, any single stride that overshoots the permissible ceiling undergoes subdivision into a sequence of smaller increments.

For each synthetic motion trace, waypoint coordinates are first stochastically generated, then interpolated into a seamless curve, and finally discretized to match the agent's per-timestep advance granularity. Displacement size L_i follows a Lévy distribution subject to an upper truncation,

$$L_i = \min \left(\xi^{-\frac{1}{\alpha}}, L_{\max} \right), \quad \xi \sim \mathcal{U}(0, 1), \quad (4)$$

Where the exponent α tunes tail mass, and $L_{\max}$ caps the permissible inter-waypoint separation.

Heading orientation advances via a Markov recurrence:

$$\theta_i = (\theta_{i-1} + \Delta\theta) \bmod 2\pi, \quad \Delta\theta \sim \mathcal{U}(-\pi S, \pi S), \quad (5)$$

With S acting as the curvature intensity coefficient, regulating the sharpness of directional shifts available to the agent.

Possessing L_i and θ_i , the refreshed coordinate pair emerges from:

$$x_i = x_{i-1} + L_i \cos \theta_i, \quad (6)$$

$$y_i = y_{i-1} + L_i \sin \theta_i, \quad (7)$$

Trace origination points are scattered randomly across the domain. Oversized jumps undergo fragmentation to remain within the ceiling speed. Relative to textbook random meanderings, the Lévy flight prescription yields motion dynamics far closer to field-observed target behavior, constituting a substantially sterner test of swarm coordination capability.

Communication model

Every drone operates a wireless return channel to a GBS, conveying its operational state and accepting synchronized maneuvering instructions in return. We cast this link as a discrete-time sequence aligned with the control clock. During any given interval t , drone i attempts delivery of a state datagram to the GBS; whether the transmission clears or gets discarded hinges on ambient network factors including spectrum contention and queue overflow.

Channel saturation and frame collisions are reproduced within the NS-3 test harness, leveraging empirically grounded radio environment models, which supply the analytical footing for the traffic management regimen introduced further below.

Baseline expert policy: GKA

Embedded within the CSOOC architecture, the Greedy Knapsack Algorithm (GKA) plays a simultaneous pair of roles: it generates oracle-grade teaching signals for the Teacher Model throughout STLF pretraining and stands as the theoretical performance ceiling during subsequent benchmarking.

GKA strips the CMOMMT formulation to a simpler premise by granting itself panoramic knowledge of every target coordinate. At each temporal slice, it performs a bipartite procedure: (1) Point Selection—opportunistically identifying vantage positions that yield the greatest immediate detection count; and (2) Assignment—applying the Hungarian matching method to shrink cumulative drone repositioning overhead. Through wedding opportunistic detection gain with mathematically optimal dispatching, GKA yields nearly ideal per-cycle moves that constitute the coordination upper bound under noise-free networking assumptions. These oracle-grade samples serve as supervisory targets for the Teacher Model within STLF, allowing it to internalize the optimal orchestration behavior achievable when complete informational access exists.

Materials and Methods

Overview

The CSOOC framework put forth herein furnishes resilient multi-drone orchestration beneath shifting and untrustworthy link conditions by tightly interlocking collective governance with live network diagnostics. Its foundational purpose is supplying uninterrupted, elastic tracking of numerous targets via connectivity-sensitive action selection. CSOOC weaves together three essential pillars:

- **Online–Offline Coordination Strategy:** Drones fluidly oscillate between centralized direction and self-sufficient operation contingent on link state. During intervals of steady GBS connectivity, drones enact commands streamed from the STLF engine (online regime); when channel saturation or frame erasure manifests, they invoke a utility-maximizing distributed policy (offline regime) to preserve sensing footprint.
- **STAR Congestion Control Mechanism:** A double-layered transmission arbitration architecture fusing chronologically desynchronized scheduling with RTS/CTS-mediated channel acquisition to suppress frame collisions and reinforce uplink steadfastness.
- **Python–NS-3 Co-Simulation Platform:** A socket-interfaced bidirectional conduit bridging the STLF inference core with the NS-3 networking sandbox, facilitating time-step-synchronized measurement of both control quality and link performance throughout each simulation cycle.
- Through merging behavioral regime switching, saturation mitigation, and cross-layer simulation, CSOOC narrows the gap between idealized algorithmic orchestration and the messy realities of wireless networking, presenting a template for cooperative multi-drone surveillance within communication-deprived IoT deployments.

Online and offline actions

The introduced CSOOC framework empowers drone collectives to alter their decision modality responsively in accordance with instantaneous link quality. When the GBS uplink holds steady, the swarm engages online mode, executing centrally synchronized directives produced by a learning-based global controller. Upon encountering link degradation or outright disruption, drones revert to offline mode, depending on locally sensed information and independent decision routines. This dual-phase operational construct guarantees dependable, uninterrupted swarm function spanning both clear-channel and impaired-connectivity circumstances.

Online mode: STLF-based centralized control

During periods of successful GBS uplink, drones function in online mode, performing movement instructions originating from the STLF system [9]. The STLF pipeline internalizes near-optimal cooperative tactics by first schooling a Teacher Model on privileged full-state inputs, then distilling this competence into a Student Model attuned to limited observability. In this way, drones can sustain harmonized, high-yield target tracking while operating beneath genuine sensor constraints.

Two-stage learning paradigm: As depicted in **Figure 4**, the STLF pipeline proceeds through two successive phases:




Figure 4. Two-stage STLF framework: Teacher Model training under complete information (Stage 1) and Student Model deployment under partial observability (Stage 2).

Stage 1–teacher model: conditioned on oracle trajectories supplied by the GKA, enjoying unconstrained global awareness. It acquires the capability to emit near-optimal displacement choices that drive the mean observation metric maximally upward.

Stage 2–student model: warm-started from the Teacher Model's parameters and subsequently honed under fragmented observability. All-encompassing global inputs are supplanted by an Observable Value Map (OVM), which yields probabilistic reconstructions of where undetected targets might reside.

Gaussian-based state representation: At every discrete time increment, the operational theater is depicted via four spatial fields that register drone conditions and target-linked data as continuous Gaussian spreads. In contrast to discrete binary occupancy arrays, Gaussian formulations supply smooth and differentiable spatial signatures, encapsulating locational ambiguity and bolstering the model's geometric inference faculties.

In particular, as portrayed in **Figure 5**, the neural architecture ingests:




Figure 5. Neural network input–output structure. Four Gaussian-encoded input maps are aggregated and processed by the STLF model to produce an action map for UAV swarm control.

UAV position map: captures the live coordinates of all drones.

The following value map reflects targets presently held inside any drone's sensory envelope, rendered as Gaussian peaks centered upon each identified target location.

Unobserved target map (teacher only) conveys the omniscient ground-truth placements of targets situated outside every drone's sensory perimeter during the Teacher Model's training phase.

Observable value map (student only) delivers a stochastic reconstruction of undetected targets, pieced together from archived sighting records throughout Student Model rollout.

These four layers are consolidated as a multi-channel input block and routed into the STLF engine, allowing concurrent deliberation over target shadowing and territory scouting under contrasting informational regimes.

Centralized action generation: The consolidated input block passes through the residual CNN backbone of the STLF pipeline, comprising successive compression, residual, and upscaling stages. The emergent output takes the form of an action map prescribing the upcoming motion selection for every drone. Such directives originate at the GBS, which

amasses status telemetry from the full drone fleet and scatters the resultant steering instructions back across the swarm. This centralized command paradigm facilitates globally harmonized cooperative maneuvers while communication pathways stay intact.

Offline mode: profit-driven local control

Whenever the GBS link breaks down—owing to channel clogging, datagram erasure, or excessive propagation lag—drones independently revert to offline functioning. Within this modality, drones no longer obtain globally synthesized STLF directives and are compelled to formulate entirely self-sufficient mobility choices. The principal purpose of the offline tactic is to uphold productive target monitoring and terrain reconnaissance notwithstanding the forfeiture of overarching coordination. This backup routine supplements the online modality inside CSOOC, guaranteeing task continuity through impaired connectivity windows.

Local map construction

Motivated by the structural conventions of the STLF input format, every drone continually manages two private, Gaussian-modeled spatial layers derived wholly from its own cumulative observational log: the Local Following Value Map (F_{local}) alongside the Local Observable Value Map (O_{local}). These layers function as stripped-down substitutes for the full-scale STLF input, allowing drones to forge autonomous choices when connectivity is absent.

Local Following Value Map (F_{local}): This layer encodes the whereabouts of targets currently caught inside the drone's detection radius. Every spotted target situated at (μ_x) is expressed through a circularly symmetric 2D Gaussian kernel.

$$f(x, y) = \frac{1}{2\pi\sigma^2} \exp \left[-\frac{(x - \mu_x)^2 + (y - \mu_y)^2}{2\sigma^2} \right], \quad (8)$$

Where σ corresponds to the drone's sensing footprint span. Multiple concurrent detections are additively overlaid, whereas the lack of any target yields a blank map, redirecting the drone's inclination toward exploration.

Local Observable Value Map (O_{local}): This layer supplies a likelihood-weighted projection regarding where unseen targets stand the best chance of surfacing, anchored exclusively on the drone's own observational timeline. It combines a temporal factor—encoding escalating uncertainty as duration since the prior visitation extends—and a spatial factor—tilting exploratory preference toward regions just outside the drone's existing sensory perimeter. The variable σ_o constrains the spatial extent of this exploratory lean.

Both layers undergo localized refreshing each timestep: F_{local} reproduces the freshest detection output, while O_{local} steadily aggregates temporal fading and spatial predilection data. This scheme lets each drone sustain a slender yet illuminating digest of its neighborhood, forming the substrate for profit-guided offline decision-making. Though the construct mirrors the global STLF blueprint in principle (Figure 5), the local layers encode solely the state data witnessed by an individual drone, eschewing any amalgamation of collective swarm-wide sightings.

Profit function and movement selection

During each offline timestep, the drone examines all feasible postures $p \in P_{\text{candidate}}$ falling inside its locomotion envelope. For any given posture, the anticipated profit is computed as the net increment in local map readings attainable upon shifting to p , gauged relative to the map generated at the drone's existing posture p_{current} :

$$\text{Profit}(p) = \sum_{x,y} [M_{\text{local}}(x, y|p) - M_{\text{local}}(x, y|p_{\text{current}})], \quad (9)$$

Where M_{local} stands for the merged local following and observable value fields.

The drone subsequently adopts the relocation that pushes this profit figure to its peak:

$$p_{\text{next}} = \arg \max_{p \in P_{\text{candidate}}} \text{Profit}(p). \quad (10)$$

This formulation conducts an explicit side-by-side weighing of the candidate-location map against the map constructed at the present position, thereby assuring that the drone settles on the posture affording the largest net uptick in informational yield.

Behavioral consistency and advantages

This utility-guided paradigm empowers drones to autonomously uphold competent coverage throughout communication gaps by marshaling localized spatiotemporal intelligence. Because F_{local} and O_{local} adhere to the identical Gaussian structure underpinning the online STLF inputs, offline comportment stays behaviorally aligned with online functioning, facilitating seamless policy reversal once connectivity resumes. This arrangement guarantees that CSOOC can preserve swarm-level operational integrity even amid pronounced communication decay.

Although CSOOC omits a dedicated collision-avoidance apparatus, it secures organic spacing via its utility-field architecture. In the online phase, targets spotted by any drone are excised from the Unobserved Target Map, diluting that locale's attractiveness and unobtrusively deterring multiple drones from gravitating toward identical coordinates. During connectivity interruption, the offline rule follows individually computed gain maps, within which formerly surveyed or low-payoff zones hold minimal lure, fostering additional geographic dispersion.

STAR: Stratified transmission and RTS/CTS congestion control

The orchestration scheme put forward in this work depends critically on swift response loops closed through the GBS. Yet a formidable difficulty stems from the inherently lockstep operating rhythm of the drone collective: as each discrete interval commences, every airborne node attempts to push its situational update toward the terrestrial receiver at the same moment. This simultaneous flood of transmissions gives rise to acute channel oversaturation, signal pile-ups, and, in the worst case, wholesale communication dropouts that eat away at aggregate surveillance outcomes.

To neutralize this bottleneck, we advance a dual-layered anti-congestion framework christened STAR. STAR fuses chronological staggering of drone emissions with an RTS/CTS-anchored medium reservation handshake, delivering hardened return-link integrity under densely packed swarm topologies.

Stratified transmission timing

Instead of condoning unison transmission from all drones within a single tick, we deploy a randomized tiered-dispersion approach that spreads emission attempts throughout each one-second block. More precisely, the time axis is carved into 1-s segments (bins) demarcated as $[k, k+1)$, with $k \in \mathbb{Z} \geq 0$ acting as the integer interval counter.

For an arbitrary drone i drawn from the set $\{1, 2, \dots, N\}$, the instant of transmission falling inside the k -th bin is dictated by:

$$t_i^{(k)} = k + U_{i,k}, \quad U_{i,k} \sim \mathcal{U}(0, 1), \quad (11)$$

Within which U_i , denotes an uncorrelated random draw governed by a uniform law over the unit interval $[0, 1]$. Every outcome in $[0, 1]$ shares an identical probability of being chosen, injecting unbiased stochasticity without privileging any

particular sub-moment of the bin.

This construct enforces exactly one datagram emission per drone per bin, whereas the randomized displacement U_i scatters emission instants across the swarm population. Through deliberate avoidance of schedule phase-locking, repeat-collision probabilities experience a meaningful drop. From an implementation standpoint, each drone runs a localized countdown that initiates datagram dispatch at the randomly drawn epoch (k) belonging to each bin.

RTS/CTS-Enhanced handshake

To place an additional barrier against data-frame crashes throughout the uplink phase, STAR weaves a Request-to-Send/Clear-to-Send (RTS/CTS [14]) dialog into the protocol machinery. Before launching its payload-bearing frame, a drone first emits a compact RTS probe aimed at the ground receiver. The ground side, upon registering the RTS request, answers with a CTS grant provided the wireless medium registers as vacant, effectively ring-fencing the channel for that particular drone.

The success or collapse of the dialog for drone i at emission opportunity ti can be characterized as:

$$S_i = \begin{cases} 1, & \text{if UAV } i \text{ receives a CTS reply from the GBS,} \\ 0, & \text{otherwise,} \end{cases} \quad (12)$$

Where the Boolean variable S_i encodes whether the RTS/CTS exchange concluded favorably.

Adopting this dialog routine alleviates the hidden-station dilemma and shrinks exposure to payload-frame pile-ups, which levy a steeper toll owing to their expanded length. The marriage of temporally de-synchronized emission slots with proactive channel reservation delivers a pronounced uplift in return-path robustness and effective capacity. The STAR regimen is embedded directly within the NS-3 emulation environment, where the medium-access sublayer faithfully instantiates both the scheduling offsets and the reservation dialog.

Beyond pure collision abatement, STAR tackles capacity ceilings and latency variability inherent in populous drone assemblies. Through forced staggering of emission windows, STAR precludes large contingents of drones from wrestling over the shared spectrum at an identical instant, consequently softening peak congestion loads.

NS-3 with python's communication

To mirror the footprint of networking vagaries upon collective orchestration fidelity, we construct a bidirectional information pipeline tethering the Python-driven CMOMMT scene engine to the NS-3 networking sandbox. Integration materializes through a socket-mediated bridge, affording live bidirectional data throughput between the two emulation domains. The communication cadence adheres to a discrete-time protocol, each interval marker t corresponding with a synchronized emulation cycle across both platforms. **Figure 6** lays out the interaction choreography per discrete moment.




Figure 6. Conceptual overview of the proposed CSOOC framework, integrating STLF-based centralized control, profit-driven offline autonomy, and STAR congestion control within a unified communication–control architecture.

As each interval opens, the Python host pushes the freshly updated positional snapshot across a TCP socket to the NS-3 engine. This payload carries the spatial coordinates of the entire drone roster, which NS-3 employs as the basis for emulating over-the-air transmissions between aerial agents and the GBS. The NS-3 kernel subsequently steps through a faithful reproduction of the uplink exchange, absorbing radio propagation decay, cross-signal interference, frame-on-frame collisions, and queuing latencies in accordance with the configured network parameterization. The yield of this emulation—namely, a per-drone binary indicator reflecting transmission outcome—gets channeled back to the Python side through the same socket pathway.

On the Python end, the CMOMMT scene handler parses the returned link-fidelity feedback to assign each drone its operational modality. Drones for which a steering directive is issued by the GBS successfully transiting the reverse path within the given interval remain under online governance, enacting the maneuver selection produced by the STLF inference pipeline. Conversely, drones for which the link proves inadequate—owing to frame erasures, buffer saturation, or untenable round-trip timing—flip into offline governance, defaulting to a self-contained locomotion rule constructed atop a privately held Gaussian gain field. This field merges discovery-driving and tracking-reinforcing incentives, equipping the drone to press forward with mission-essential duties independently of any overarching coordinator.

The socket-based coupling maintains strict temporal lockstep between both simulation realms over the entire mission timeline. Each Python-hosted advancement step translates to a constant-duration slice inside NS-3, upholding chronometric congruence between the movement stratum and the networking stratum. This joint simulation facility opens the door to probing the online policy's resilience across a spectrum of link-quality regimes, furnishing a more authentic reading of swarm aptitude within communication-compromised field conditions.

Experiment

Experimental setup

Empirical runs were executed under Python 3.9 leveraging the PyTorch 2.9.1 neural computation suite, with training conducted on an NVIDIA V100 accelerator. The networking emulation subsystem was housed within a VMware® Workstation 16 Pro virtualized container (version 16.2.3) operating Ubuntu 22.04.2 LTS. The NS-3 networking emulator (version 3.44) served to recreate wireless propagation conditions, complete with saturation episodes, datagram loss, and transit delay artifacts, drawing upon its native Wi-Fi and LTE protocol stacks. The Python-resident STLF environment and the NS-3 instance were synchronized through TCP sockets, preserving unwavering discrete-time messaging and actuation loops across every simulation trajectory.

Table 1 enumerates the simulation settings selected to realize the CMOMMT scenario. The survey arena takes the form of a 400×400 cellular grid furnished with 60 mobile subjects. Drone fleet magnitudes are instantiated at 15, 20, and 30 units, translating to drone-per-subject quotients of 1:4, 1:3, and 1:2, respectively. Sensing envelopes are provisioned at radii of 12 (restricted), 18 (intermediate), and 24 (wide) cells. Drones and subjects traverse at top speeds of 3 and 1 cells per interval, respectively. Every run unfolds over 500 intervals. Reported metrics are averaged over five uncorrelated repetitions to suppress stochastic fluctuations traceable to random initial seeding and transmission-time jitter.

Table 1. Simulation Parameters.

Parameter	Value
Search Region	400×400 Cells
Number of Targets	60
Number of UAVs	15/20/30
UAV Observation Radius	12/18/24 Cells
UAV Maximum Speed	3 Cells per Step
Target Maximum Speed	1 Cell per Step
Mission Execution Period	500 Time Steps

Inside the NS-3 parameterization, airborne nodes converse with the terrestrial base station via Wi-Fi (IEEE 802.11n) at a constant emission strength of 10 dBm, spanning a 20 MHz channel. Every drone issues a lone UDP datagram sized at 1024 Bytes per interval, enabling high-fidelity reproduction of overload-triggered frame shedding amid dense node concentrations.

With respect to the Gaussian blurring enacted upon input feature layers, the spread coefficient σ is tuned in a layer-specific fashion: within the drone positional layer, σ is matched to the drone's ceiling displacement span per interval; inside the following value layer, σ aligns with the drone's detection footprint radius; and within the unseen subject location layer, σ absorbs a magnitude equal to one-and-a-half times the drone's detection footprint radius. This differentiated calibration ensures each feature layer conveys its designated spatial footprint and sensing ambiguity with high fidelity.

Evaluation metrics

To render a numerical assessment of CSOOC's effectiveness, we draw upon two distinct families of measures:

(1) Mission-level measures (CMOMMT attainment): The foremost indicator of success is the Average Observation Rate (AOR), computed as the quotient formed by dividing the total tally of targets spotted at each interval by the overall number of targets present within the operational zone. A loftier AOR denotes enhanced fleet-wide spatial dominion and tracking throughput sustained over the entire sortie.

(2) Link-level measures (transmission fidelity): To delineate the imprint of networking conditions upon collective coordination, we record:

Packet Delivery Ratio (PDR): the proportion of datagrams that land intact at the GBS, measured against the complete set of datagrams dispatched.

Latency: the averaged door-to-door propagation lag incurred by packets journeying from aerial platforms to the terrestrial receiver.

Jitter: the degree of inconsistency in packet inter-arrival gaps, functioning as an index of wireless pathway steadiness.

Collectively, these gauges span both operational accomplishment and the foundational networking robustness, affording an encompassing appraisal of the submitted architecture beneath disparate connectivity scenarios.

Validation with NS-3 network simulation

To lay bare the performance attrition inflicted by transmission degradations, a deliberately restrained-scale setup is enlisted, featuring 15 drones each possessing a detection footprint spanning 12 units. **Figure 7** charts the averaged observation rate progressions under this specification. In the interest of completeness, **Figure 7** additionally depicts the variability envelope (spanning ± 1 standard deviation) hugging each trajectory as a translucent band. The SD quantities register at 0.76% for GKA coupled with NS-3, 0.78% for GKA decoupled from NS-3, 1.56% for STLF coupled with NS-3, and 1.16% for STLF decoupled from NS-3.




Figure 7. Average observation rate under ideal and NS-3 simulated network conditions.

Examining the curves that include NS-3, one observes that both STLF and GKA uphold consistent pursuit behavior, albeit a readily apparent capability chasm endures. GKA, privileged with exhaustive global situational insight, unfailingly crafts near-optimal dispatch selections, translating to a superior observation fraction. STLF, notwithstanding its capacity to mimic the oracle regimen through knowledge transplantation, yields somewhat curtailed returns owing to the strictures of fragmentary awareness.

When the NS-3 networking stratum is interlaced into the operational pipeline, overall capability contracts. Since the entirety of the drone contingent seeks to push status payloads upward at the very same instant marking each interval boundary, channel saturation, frame clashes, and supplementary buffering holdups manifest. These corruptions strike both techniques with pronounced severity, as fickle connectivity undercuts the exactness of synchronized action selection, thereby depressing the averaged observation fraction relative to the noiseless baseline. This experimental run additionally substantiates that our socket-bridged co-simulation conduit faithfully unites the Python-resident CMOMMT domain with the NS-3 networking testbed. The measured slump in both transmission fidelity and observational return squares entirely with analytical expectations, verifying the robustness and functional integrity of the coupling.

Table 2 itemizes the networking diagnostics. One may note that PDR remains beneath full completion for both schemes, registering at 75.35% for GKA and 71.91% for STLF. In parallel, averaged latency dwells in the neighborhood of 5 ms, while jitter figures land at 3.26 ms and 3.61 ms, respectively. These readings intimate that although the pathway stays partially viable, datagram disappearance and temporal dispersion levy a straight erosive toll upon centralized orchestration, thus compressing aggregate pursuit yield.

Table 2. NS-3 network metrics for GKA and STLF.

Method	STLF	GKA
PDR (%)	71.91	75.35
Latency (ms)	5.38	5.33
Jitter (ms)	3.61	3.26

To encapsulate, **Figure 7** together with **Table 2** jointly substantiate that: (1) operating beneath flawless connectivity, GKA reliably surpasses STLF; (2) after injecting genuine networking artifacts, both methodologies register capability diminishment; and (3) the presented Python–NS-3 socket-interlinked co-simulation substrate accurately reproduces network-traceable performance erosion, confirming its value for communication-sensitive drone fleet coordination inquiries.

Pursuing broader NS-3 corroboration, we interrogate the repercussions of modulating channel bandwidth, benchmarking both methodologies beneath 10 MHz alongside 40 MHz provisions while preserving 20 MHz as the reference setup. The 20 MHz allocation adheres to the standard IEEE 802.11 PHY parameterization employed uniformly within our NS-3 experiments and delivers ample headroom for the experimental traffic burden. As tabulated in **Table 3**, capability persists near-constant when the allocation is broadened to 40 MHz, certifying that the reference setting already satisfies traffic requirements without strain. Oppositely, narrowing the allocation to 10 MHz triggers observable deterioration in both PDR and AOR, signposting that the spectrum grows saturated once channel capacity is curtailed.

Table 3. Performance under different bandwidth settings.

Metrics	10 MHz	20 MHz	40 MHz
GKA			
AOR (%)	17.83	20.19	20.03
PDR (%)	63.87	75.35	74.40
Latency (ms)	14.15	5.33	1.14
Jitter (ms)	11.03	3.26	3.42
STLF			
AOR (%)	10.73	13.96	14.21
PDR (%)	63.90	71.91	73.33
Latency (ms)	13.00	5.38	3.62
Jitter (ms)	10.13	3.61	2.82

Validation of STAR congestion control

To size up the effectiveness of the submitted STAR construct, we place networking performance and fleet-level pursuit statistics side by side under configurations where STAR is operative against configurations where it is withheld, preserving all remaining transmission and operational parameters unchanged. In practice, experimental runs are performed adopting the 15-drone, 12-unit detection radius arrangement, scrutinizing both link-layer statistics and swarm-scale observation measurements.

The figures assembled in **Table 4** transparently attest that engaging the STAR congestion-dampening apparatus procures sizeable enhancements in pathway dependability and collective harmonization output. More specifically, PDR ascends from 75.35% to 92.96% in the GKA instance and from 71.91% to 94.22% in the STLF instance, denoting a substantial shrinkage in frame clashes and delivery misfires. With pathway consistency strengthening, the drone fleet accesses more trustworthy orchestration commands, giving rise to augmented target envelopment and following precision.

Table 4. Performance comparison with and without STAR congestion control.

Metrics	With STAR	Without STAR
GKA		
AOR (%)	22.99	20.19
PDR (%)	92.96	75.35
Latency (ms)	5.82	5.33
Jitter (ms)	1.88	3.26
STLF		
AOR (%)	17.44	13.96
PDR (%)	94.22	71.91
Latency (ms)	5.85	5.38
Jitter (ms)	3.52	3.61

Beyond this, STAR regularizes the pathway by measurably suppressing jitter. As an illustration, the jitter metric for GKA shrinks from 3.26 ms to 1.88 ms, suggestive of heightened predictability and diminished burst-like emission lags. Albeit averaged latency registers a mild upward tick for both schemes (moving, for example, from 5.33 ms to 5.82 ms in the GKA case), this trade-off remains justifiable when set against the considerable advances in delivery certainty and temporal smoothness. The extra queuing time traces back to STAR's temporal layering element, which purposefully spaces out drone emissions to alleviate pile-ups.

Seen as a whole, these empirical patterns substantiate the potency of the submitted STAR congestion-dampening apparatus. Through the combined application of chronological de-synchronization in concert with anticipatory channel booking, STAR consolidates pathway steadiness and lifts fleet harmonization capability within high-density transmission contexts.

Validation of offline action within CSOOC

To single out and weigh the incremental merit of offline maneuvering housed inside the CSOOC framework, we orchestrate controlled trials in the presence of connectivity unpredictability. These trials reuse the identical parameter specification employed in the networking corroboration experiment—15 drones provisioned with a detection footprint spanning 12 units.

When GBS-originated communication feedback vanishes, drones operating under CSOOC pivot from online functioning toward the offline maneuvering subsystem, navigated by Gaussian-shaped localized gain landscapes. Rather than becoming inert or defaulting to aimless roaming, each drone independently formulates its locomotion selections anchored in its privately stored observational chronicle, fluidly arbitrating between target pursuit and reconnaissance sweeping. This formulation ensures behavioral uninterruptedness within the vacuum produced by absent centralized direction, upholding competent surveillance over mobile entities.

As presented in Table 5, incorporating offline maneuvering generates a measurable advancement in pursuit figures. Regarding the GKA baseline, AOR progresses from 20.19% to 23.46%, whereas for STLF, the value rises from 13.96% to 17.78%. PDR experiences a parallel uplift, mirroring the heightened adaptiveness of drone conduct throughout stretches of volatile linkage. Figure 8 further tracks how the averaged observation fraction unfolds along the mission horizon, revealing upward trends for both GKA and STLF once offline maneuvering is brought into play.




Figure 8. Average observation rate evolution with and without STAR.

Table 5. Performance comparison with and without offline action.

Metrics	With offline action	Without offline action
GKA		
AOR (%)	23.46	20.19
PDR (%)	82.55	75.35
Latency (ms)	6.36	5.33
Jitter (ms)	3.96	3.26
STLF		
AOR (%)	17.78	13.96
PDR (%)	80.96	71.91

Latency (ms)	6.35	5.38
Jitter (ms)	3.99	3.61

These findings throw into relief the essential function shouldered by the offline maneuvering module within CSOOC:

- It preserves harmonized fleet deportment by way of distributed decision-making, even while global steering is momentarily beyond reach.
- It assures graduated performance tapering in place of abrupt operational collapse during linkage interruptions.
- It furnishes a connective tissue spanning online and offline modalities, preserving tactical alignment with the STLf policy.

In summation, the offline module operates as an indispensable behavioral chassis of the CSOOC architecture, endowing drone collectives with the capacity to sustain hardy observational throughput amid connectivity unpredictability absent any requirement for external remediation.

Validation of the CSOOC and comparative analysis

So as to undertake a thorough corroboration of the submitted CSOOC architecture's merit, we set its output side by side with that of a suite of reference algorithms, varying both the ceiling speed permitted to drones and the numerical balance between drones and targets. In the foregoing experimental sections, the separate value added by the offline maneuvering element and by the STAR congestion dampening element was scrutinized individually; now, for the present comparisons, both building blocks are simultaneously activated inside every method being benchmarked. This choice guarantees a level playing ground and permits the influence attributable strictly to coordination logic to be isolated. Aiming to further test how well CSOOC generalizes, the evaluation spans multiple collective dimensions (drone counts of 15, 20, and 30) alongside several sensing envelopes (radii of 12, 18, and 24 cells). These variations effectively modify the drone-to-target proportion and the degree of spatial overreach within the CMOMMT testbed. The collection of comparator algorithms encompasses:

- GKA: A near-optimal rule-based procedure that demarcates the theoretical achievement ceiling.
- CSOOC: A fortified outgrowth of STLf that folds in both a utility-weighted locomotion routine and congestion suppression, yielding connectivity-adaptive collective orchestration.
- PAMTS [5]: PAMTS is a gain-driven seek-and-tail scheme that aligns drones by means of an observational payoff measure, routing each aerial unit at every tick toward the spatial cell carrying the loftiest estimated reward.
- PAMTS Follow [5]: A branch of PAMTS that tips the balance toward tailing once a target has been acquired. Standing apart from the base PAMTS variant—which continuously negotiates the trade-off between fresh terrain search and ongoing pursuit—PAMTS Follow entirely silences the search drive by forcing its weighting term to zero. This yields steadier observational continuity, a trait especially valuable in assignments that call for unbroken vigilance, typified by convoy guarding or extended-duration watch.
- A-CMOMMT [4]: The chronologically first documented attempt at resolving the CMOMMT formulation.
- I-CMOMMT [7]: Layered over A-CMOMMT, I-CMOMMT brings in a vector encoding how long each zone has gone unvisited since the last sighting sweep.
- ROD (Reward-based Orientation Determination) [8]: The ROD technique is placed within the comparison pool owing to its approach of settling heading choices via reward scoring to fine-tune the efficiency of joint target location.
- VDP-ACO (Voronoi-based Dynamic Partition with Ant Colony Optimization) [33]: A decentralized construct that unites Voronoi-driven adaptive area carving with ant colony metaheuristics, operating beneath the umbrella of distributed model predictive control to boost collaborative target location performance inside shifting surroundings.

The comparator suite assembled for this study collects emblematic entries drawn from the CMOMMT literature, cutting across traditional, optimization-built, and learning-infused coordination doctrines that engage head-on with the distinctive traits of the multi-target shadowing task. While it is true that DRL-based coordination algorithms also warrant attention, the emphasis of the present evaluation falls on positioning CSOOC opposite approaches tailored specifically to the problem domain and to the task at hand. Absorbing DRL-founded coordination paradigms stands as a significant avenue to be pursued in follow-on investigations.

As rendered visible through **Figure 9** and **Table 6**, the mean observation percentage yielded by every method rises in steady lockstep with increments in both drone count and sensing footprint radius. This pattern matches expectation, considering that additional aerial nodes and broader sensing swaths jointly translate into more expansive territorial command and heightened rates of target interception. Across the entirety of the parameter landscape, the performance differentials separating the methods incrementally shrink as the collective grows larger, a phenomenon that captures the waning marginal return of orchestration effort once near-complete coverage has been approached. Still, unmistakable performance rankings persist among the compared methods.




Figure 9. Average observation rate comparison under different UAV counts and observation radii.

Table 6. Average Observation Rate (%) for different UAV numbers and observation radii.

r	Algorithm	15 UAV	20 UAV	30 UAV
12	GKA	28.23 ± 1.21	35.50 ± 2.11	53.47 ± 2.84
	CSOOC	22.47 ± 0.73	25.15 ± 1.42	42.57 ± 2.05
	PAMTS	3.93 ± 0.02	9.26 ± 0.13	38.58 ± 1.85
	PAMTS Follow	14.50 ± 1.39	19.95 ± 2.23	36.99 ± 1.96
	A-CMOMMT	18.83 ± 0.91	22.24 ± 1.64	31.69 ± 1.47
	I-CMOMMT	17.42 ± 0.66	20.80 ± 1.15	32.32 ± 1.99
	ROD	10.98 ± 0.31	16.48 ± 0.53	28.29 ± 1.58
	VDP-ACO	16.98 ± 1.57	19.96 ± 0.86	24.82 ± 1.39
18	GKA	40.15 ± 1.27	47.70 ± 1.93	62.40 ± 2.88
	CSOOC	29.63 ± 0.74	36.91 ± 1.12	49.42 ± 2.41
	PAMTS	15.73 ± 1.03	21.64 ± 2.08	47.56 ± 3.08
	PAMTS Follow	28.50 ± 1.44	35.08 ± 1.34	42.34 ± 3.19
	A-CMOMMT	28.58 ± 0.96	33.16 ± 2.31	44.82 ± 2.57
	I-CMOMMT	26.62 ± 1.51	34.08 ± 1.27	46.38 ± 2.02
	ROD	21.38 ± 0.85	29.12 ± 2.46	38.69 ± 2.91
	VDP-ACO	23.36 ± 1.63	23.96 ± 1.18	39.95 ± 1.99
24	GKA	51.17 ± 2.21	61.84 ± 2.16	75.99 ± 3.12
	CSOOC	39.84 ± 1.08	47.50 ± 1.57	63.87 ± 2.65
	PAMTS	29.55 ± 1.93	36.21 ± 2.44	56.24 ± 1.88
	PAMTS Follow	38.61 ± 1.36	41.80 ± 1.73	52.68 ± 2.07
	A-CMOMMT	36.71 ± 1.77	44.44 ± 2.03	53.76 ± 2.49
	I-CMOMMT	35.60 ± 1.12	46.22 ± 1.82	58.29 ± 3.31
	ROD	29.16 ± 1.25	34.35 ± 1.66	48.70 ± 1.93

The put-forward CSOOC method delivers steady and elevated performance throughout every configuration submitted to testing. Through the fusion of a utility-weighted offline component with the STAR congestion dampening module, CSOOC dependably sustains elevated observation scores under conditions of both low and high drone density. Weighed against the pre-improvement STLF, CSOOC shows superior elasticity in the face of shifting drone-to-target ratios, managing to keep coordination intact even during spells of communication intermittency. Its uniform uplift across every tested radius value throws into relief the architecture's knack for smoothly meshing top-down and bottom-up decision modes while navigating connection unreliability.

By comparison, time-honored heuristic approaches exhibit comparatively brittle behavior. PAMTS and A-CMOMMT rest heavily on parameter hand-tuning and hardwired rule sets, factors that clip their capacity to adjust when the arena scale shifts. PAMTS Follow and ROD each carry a pronounced behavioral tilt—toward tailing in the former case, toward

scouting in the latter—giving rise to an unsatisfactory compromise between near-continuous observation and far-reaching spatial canvassing. I-CMOMMT reaps incremental headway through its awareness of recent visitation lapse, yet stays hemmed in by the absence of any communication feedback model. VDP-ACO produces oscillating outcomes: it puts up competitive numbers within mid-sized operational settings but sheds efficacy inside densely populated configurations, a penalty exacted by its distributed optimization burden.

When reduced to a single aggregate figure, CSOOC clocks a mean observation fraction of 39.7% over the entire sweep of parameter combinations, clearing the competing baseline algorithms by spans of 4.4%–11.13%. These data points confirm that CSOOC both preserves stable orchestration beneath communication bottlenecks and transfers effectively across heterogeneous collective sizes and sensing footprints, cementing its standing as a rugged and broadly applicable answer to real-deployment CMOMMT problems.

The network-layer performance statistics assembled in **Table 7** supply additional weight to the demonstrated potency of the STAR congestion-dampening module embedded uniformly across all tested methods. Placed alongside the matching readings from **Table 4** (where STAR was withheld), every algorithm posts sizeable boosts in datagram delivery fraction (PDR), together with compressed transit lag and dampened temporal variability, thereby affirming that STAR succeeds in curbing medium contention and bringing stability to channel behavior. Mean PDR figures across the full configuration matrix rise by increments of 13.68–21.48%, whereas both lag and jitter retreat, indicative of more faithful and smoother return-path exchange bridging the aerial nodes and the terrestrial terminal. Rather than laboriously stepping through each cell of **Table 7**, we draw out the overarching motifs that capture the collective network profile. The STAR module consistently clamps the datagram delivery fraction at elevated levels, routinely exceeding the 85–90% band over heterogeneous collective magnitudes and detection footprints. This observation signals that STAR effectively tamps down saturation and raises communication constancy as drone numbers change. CSOOC holds end-to-end transit lag beneath the 10 ms ceiling throughout every configuration examined, a sign that the architecture preserves real-time reactivity under shifting network burdens.

Table 7. Network performance under different UAV numbers and observation radii. PDR in %, Delay/Jitter in ms.

<i>r</i>	Algorithm	15 UAV			20 UAV			30 UAV		
		PDR	Delay	Jitter	PDR	Delay	Jitter	PDR	Delay	Jitter
12	GKA	92.97	5.40	1.57	89.47	11.41	8.24	91.06	5.37	1.63
	CSOOC	93.36	6.19	2.57	88.93	10.71	7.48	86.55	6.04	2.63
	PAMTS	92.18	10.66	8.66	87.30	13.3	11.01	90.10	10.06	7.10
	PAMTS Follow	94.01	5.03	1.41	90.01	9.96	7.24	88.10	8.82	5.77
	A-CMOMMT	93.10	5.56	1.78	88.69	11.04	7.71	91.24	5.82	2.17
	I-CMOMMT	91.74	6.10	2.57	91.98	10.13	6.46	87.66	9.57	6.07
	ROD	91.14	6.49	3.12	87.89	10.68	7.05	86.58	11.27	9.37
	VDP-ACO	92.86	6.55	3.62	91.79	11.55	8.38	88.07	7.98	5.37
18	GKA	91.38	5.65	2.06	91.51	5.67	1.88	90.49	5.64	1.87
	CSOOC	90.34	6.81	3.26	90.06	6.13	2.76	87.83	11.10	7.55
	PAMTS	90.37	10.27	7.93	87.07	17.92	15.32	91.05	14.49	11.65
	PAMTS Follow	93.48	8.09	5.21	86.23	14.64	11.35	93.19	8.88	5.84

	A-CMOMMT	91.76	11.59	8.11	90.32	14.73	11.01	90.08	12.07	9.88
	I-CMOMMT	91.71	9.22	6.28	89.37	16.85	13.63	90.67	12.94	9.51
	ROD	85.17	10.76	8.42	89.25	12.89	9.96	85.30	15.46	12.41
	VDP-ACO	90.24	9.11	7.05	90.27	16.95	13.52	87.68	12.59	10.34
24	GKA	92.93	5.82	1.92	92.00	5.90	1.72	88.45	5.33	1.74
	CSOOC	90.80	6.80	3.11	91.03	8.85	5.43	89.38	9.31	5.57
	PAMTS	89.28	10.53	8.14	88.76	7.84	4.53	91.57	7.67	4.77
	PAMTS Follow	93.39	6.52	2.51	89.38	7.15	3.82	87.06	7.63	4.89
	A-CMOMMT	91.47	7.32	3.55	91.89	8.53	5.90	86.29	8.83	6.31
	I-CMOMMT	92.29	7.78	4.67	90.48	7.97	4.78	86.26	15.46	11.03
	ROD	91.36	7.68	3.92	91.75	7.45	4.33	87.35	10.58	7.06
	VDP-ACO	91.83	7.34	4.21	86.97	8.52	5.97	85.59	10.76	7.97

It deserves pointing out that the GKA and CSOOC entries do not invariably snatch the absolute top networking statistics among the full set of evaluated techniques. This spread, far from pointing to a flaw in the proposed framework, originates instead from phenomena rooted in spatial arrangement. Owing to the placement of the terrestrial terminal at the map centroid (200,200), drones that loiter near that central point enjoy steadier links, while those flung toward the periphery wrestle with fainter signal power and fiercer channel competition. The aggregate networking summary, in consequence, reflects the geographic scatter of drones and their moment-by-moment range to the terminal, rather than properties of the steering algorithm itself.

Even in the face of these variations, CSOOC sustains an unwaveringly strong PDR (north of 90%) alongside tightly bounded latency (south of 10 ms) throughout the entire roster of collective arrangements. This steadiness underwrites timely feedback flow for the online modality and trustworthy oscillation between online and offline comportment during episodes of connectivity fluctuation. Reading the results as a whole, it becomes evident that STAR makes a tangible contribution to hardening the underlying communication fabric, and that CSOOC stands capable of safeguarding durable orchestration output even as the physical geometry or link conditions shift in unpredictable ways.

Conclusion

The present contribution has put forward CSOOC, an architecture for online–offline coordination steered by communication-state awareness, targeting the problem of multi-target shadowing by drone collectives immersed in realistic networking settings. Extending the foundation laid by the STLF design, CSOOC introduces two decisive expansions: a utility-weighted offline maneuvering element for self-sufficient control across intervals of connection blackout, and the STAR congestion dampening module dedicated to preserving return-path steadiness within crowded swarm deployments. Through the real-time coupling of collective decision logic with network communication diagnostics, CSOOC attains rugged and elastic orchestration, one that remains capable of upholding elevated observational yield even when connectivity falls short of ideal.

An extensive battery of experiments, staged on a Python–NS-3 co-simulation test rig, has borne witness to the architecture's potency and capacity to generalize. The collected evidence shows CSOOC reliably outstripping baseline comparators over a wide spectrum of parameter settings encompassing drone tallies, sensing footprint radii, and target packing densities. The inclusion of STAR pushes datagram delivery fractions upward by better than 15% in averaged terms, and the offline element provides an unbroken behavioral thread across transient link dropouts. Taken in concert,

these results validate the proposition that communication-cognizant orchestration paired with responsive control can narrow the chasm separating algorithmic efficiency from the gritty realities of drone collective deployment inside actual IoT sensing contexts.

Looking past the simulation domain, CSOOC is structured to transfer smoothly into operational settings where communication quality undergoes temporal swings. Representative application arenas embrace broad-acre agricultural surveillance, emergency response sorties, persistent ecological monitoring, and extended-range infrastructure surveillance—contexts within which uplink bottlenecks and patchy connectivity emerge as commonplace challenges. The STAR mechanism offers the means to regularize datagram conveyance through such adversities, while the online–offline mode-transition logic furnishes the collective with the ability to sustain harmonized conduct even as momentary communication breaches occur. Trajectories for ensuing research will center on generalizing the CSOOC scaffold toward decentralized networking topologies and mixed-fleet UAV configurations, as well as on the incorporation of learning-powered congestion projection and dynamic spectrum allocation, steps intended to further scale up robustness and extend reach within expansive aerial meshes.

Acknowledgements

None

Conflict of interest

None

Financial support

This research was funded by the Science and Technology Development Fund of Macau, grant number 0079/2019/AMJ.

Ethics statement

None

Received: 05 Sep 2021

Revised: 29 Oct 2021

Accepted: 22 Nov 2021

Published online: 15 March 2022

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

Wei Z, Zhu M, Zhang N, Wang L, Zou Y, Meng Z, et al. UAV-assisted data collection for Internet of Things: A survey. *IEEE Internet Things J.* 2022;9(17):15460–83.

Messaoudi K, Oubbati OS, Rachedi A, Lakas A, Bendouma T, Chaib N. A survey of UAV-based data collection: Challenges, solutions and future perspectives. *J Netw Comput Appl.* 2023;216:103670.

Gupta L, Jain R, Vaszkun G. Survey of important issues in UAV communication networks. *IEEE Commun Surv Tutor.* 2016;18(2):1123-52.

Parker LE. Distributed algorithms for multi-robot observation of multiple moving targets. *Auton Robot.* 2002;12(3):231-55.

Li X, Chen J, Deng F, Li H. Profit-driven adaptive moving targets search with UAV swarms. *Sensors.* 2019;19(7):1545.

Parker LE, Emmons BA. Cooperative multi-robot observation of multiple moving targets. In: *Proceedings of the 1997 IEEE International Conference on Robotics and Automation (ICRA)*; 1997 Apr 20-25; Albuquerque, NM, USA. Vol. 3. p. 2082-9.

Chahal J, Belbachir A, Seghrouchni AEF. I-CMOMMT: A multiagent approach for patrolling and observation of mobile targets with a continuous environment representation (S). In: *Proceedings of the 33rd International Conference on Software Engineering & Knowledge Engineering (SEKE)*; 2021 Jul 1-10; Pittsburgh, PA, USA. p. 21-4.

Luo Q, Luan TH, Shi W, Fan P. Edge computing enabled energy-efficient multi-UAV cooperative target search. *IEEE Trans Veh Technol.* 2023;72(6):7757-71.

Sun H, Li X, Yan Y, Jiang T, Liu B. A supervised and transfer learning based two-stage framework for UAV swarm multi-target tracking. *J Latex CI Files.* 2025;14(8):1-15.

Nawaz H, Ali HM, Laghari AA. UAV communication networks issues: A review. *Arch Comput Methods Eng.* 2021;28(3):1349-69.

Meng K, Wu Q, Xu J, Chen W, Feng Z, Schober R, et al. UAV-enabled integrated sensing and communication: Opportunities and challenges. *IEEE Wirel Commun.* 2024;31(1):97-104.

Pourjabar M, AlKatheeri A, Rusci M, Barcis A, Niculescu V, Ferrante E, et al. Land & localize: An infrastructure-free and scalable nano-drones swarm with UWB-based localization. In: *Proceedings of the 2023 19th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT)*; 2023 Jun 19-21; Pafos, Cyprus. p. 654-60.

Nowakowski M, Idzkowski A. Ultra-wideband signal transmission according to European regulations and typical pulses. In: *Proceedings of the 2020 International Conference Mechatronic Systems and Materials (MSM)*; 2020 Jul 1-3; Białystok, Poland. p. 1-4.

Xu K, Gerla M, Bae S. Effectiveness of RTS/CTS handshake in IEEE 802.11 based ad hoc networks. *Ad Hoc Netw.* 2003;1(1):107-23.

Kolling A, Carpin S. Cooperative observation of multiple moving targets: An algorithm and its formalization. *Int J Robot Res.* 2007;26(9):935-53.

Ding Y, Zhu M, He Y, Jiang J. P-CMOMMT algorithm for the cooperative multi-robot observation of multiple moving targets. In: *Proceedings of the 2006 6th World Congress on Intelligent Control and Automation*; 2006 Jun 21-23; Dalian, China. Vol. 2. p. 9267-71.

Ding Y, He Y. Flexible formation of the multi-robot system and its application on CMOMMT problem. In: *Proceedings of the 2010 2nd International Asia Conference on Informatics in Control, Automation and Robotics (CAR)*; 2010 Mar 6-7; Wuhan, China. Vol. 1. p. 377-82.

Li B, Wang J, Song C, Yang Z, Wan K, Zhang Q. Multi-UAV roundup strategy method based on deep reinforcement learning CEL-MADDPG algorithm. *Expert Syst Appl.* 2024;245:123018.

Yan J, Zhang L, Yang X, Chen C, Guan X. Communication-aware motion planning of AUV in obstacle-dense environment: A binocular vision-based deep learning method. *IEEE Trans Intell Transp Syst.* 2023;24(12):14927-43.

Psomiadis E, Maity D, Tsiotras P. Communication-aware map compression for online path-planning. In: *Proceedings of the 2024 IEEE International Conference on Robotics and Automation (ICRA)*; 2024 May 13-17; Yokohama, Japan. p. 12368-74.

Islam N, Rashid MM, Pasandideh F, Ray B, Moore S, Kadel R. A review of applications and communication technologies for Internet of Things (IoT) and unmanned aerial vehicle (UAV) based sustainable smart farming. *Sustainability.* 2021;13(4):1821.

Eltahlawy AM, Aslan HK, Abdallah EG, Elsayed MS, Jurcut AD, Azer MA. A survey on parameters affecting MANET performance. *Electronics.* 2023;12(9):1956.

ns-3 Consortium. ns-3: Discrete-Event Network Simulator [Internet]. [cited 2025 Oct 22]. Available from: <https://www.nsnam.org/>

Piro G, Baldo N, Miozzo M. An LTE module for the ns-3 network simulator. In: *Proceedings of the 4th International ICST Conference on Simulation Tools and Techniques (SimuTools)*; 2011 Mar 22-24; Barcelona, Spain. p. 415-22.

Gill JS, Velashani MS, Wolf J, Kenney J, Manesh MR, Kaabouch N. Simulation testbeds and frameworks for UAV performance evaluation. In: *Proceedings of the 2021 IEEE International Conference on Electro Information Technology (EIT)*; 2021 May 14-15; Mt. Pleasant, MI, USA. p. 335-41.

Jiang W, Han H, He M, Gu W. Network simulation tools for unmanned aerial vehicle communications: A survey. *Int J Commun Syst.* 2024;37(10):e5878.

Baidya S, Shaikh Z, Levorato M. Flynetsim: An open source synchronized UAV network simulator based on ns-3 and ArduPilot. In: *Proceedings of the 21st ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM)*; 2018 Oct 28-Nov 2; Montreal, QC, Canada. p. 37-45.

Garg S, Ihler A, Bentley ES, Kumar S. A cross-layer, mobility, and congestion-aware routing protocol for UAV networks. *IEEE Trans Aerosp Electron Syst.* 2023;59(4):3778-96.

Feng Y, Wang Y, Zhang B, Zhou L. Deep reinforcement learning based TCP congestion control in UAV assisted wireless networks. In: *Proceedings of the 2023 International Conference on Wireless Communications and Signal Processing (WCSP)*; 2023 Nov 2-4; Hangzhou, China. p. 862-7.

Kang M, Jeon SW. Energy-efficient data aggregation and collection for multi-UAV-enabled IoT networks. *IEEE Wirel Commun Lett.* 2024;13(4):1004-8.

Hamissi A, Dhraief A. A survey on the unmanned aircraft system traffic management. *ACM Comput Surv.* 2023;56(6):1-37.

Naqvi HA, Hilman MH, Anggorojati B. Implementability improvement of deep reinforcement learning based congestion control in cellular network. *Comput Netw.* 2023;233:109874.

Li Y, Zhang Z, Sun Q, Huang Y. A distributed framework for multiple UAV cooperative target search under dynamic environment. *J Frankl Inst.* 2024;361(13):106810.