International Journal of Micro Air Vehicles

ORIGINAL RESEARCH ARTICLE

Open Access

https://doi.org/10.48130/mav-0026-0005

International Journal of Micro Air Vehicles **2026**, 18: e004

Capability-centric overlapping coalition formation for micro air vehicle swarm in time-varying communication topologies

Xin Liao[1,2], Tao Yang[1], Wenxing Fu[1], Guoda Cheng[3] and Fei Yan[4*]

[1] *Unmanned System Research Institute, National Key Laboratory of Unmanned Aerial Vehicle Technology, Integrated Research and Development Platform of Unmanned Aerial Vehicle Technology, Northwestern Polytechnical University, Xi'an 710072, China*
[2] *Shanghai Electro-Mechanical Engineering Institute, Shanghai 201109, China*
[3] *Aviation Industry Corporation of China (AVIC) Xi'an Flight Automatic Control Research Institute, Xi'an 710065, China*
[4] *School of Automation, Xi'an University of Posts and Telecommunications, Xi'an 710121, China*
* Correspondence: yanfei@xupt.edu.cn (Yan F)

Abstract

Micro air vehicle (MAV) swarm is increasingly deployed in complex cooperative search-and-execute missions, in which efficient task allocation becomes critical because of severe resource and communication constraints. Although overlapping coalition formation (OCF) has shown advantages in resource utilization compared with traditional disjoint approaches, the significant impact on time-varying communication topologies caused by the high mobility and limited communication range typical of MAV platforms is mostly neglected. To tackle these challenges in realistic MAV swarm operations, this paper proposes a capability-centric OCF method specifically tailored for dynamic, communication-constrained environments. We present the Topology-Adaptive Capability Matching (TACM) algorithm, which combines a market-based auction mechanism with capability-driven heuristics. TACM enables MAVs to dynamically aggregate heterogeneous capabilities on-the-fly while strictly adhering to local communication constraints through a hop-limited propagation strategy. Additionally, an equitable capability consumption mechanism is introduced to balance resource depletion across swarm members, thereby improving the long-term operational resilience and robustness of the MAV swarm. Results from extensive simulation demonstrate that the proposed method effectively adapts to the rapid topology changes induced by MAV mobility. Comparative studies show that TACM significantly outperforms the polynomial time coalition formation algorithm and the multi-objective quantum genetic algorithm in terms of average mission completion time and overall system utility under the time-varying communication conditions typical of MAV applications.

Citation: Liao X, Yang T, Fu W, Cheng G, Yan F. 2026. Capability-centric overlapping coalition formation for micro air vehicle swarm in time-varying communication topologies. *International Journal of Micro Air Vehicles* 18: e004 https://doi.org/10.48130/mav-0026-0005

Introduction

With the rapid miniaturization, cost reduction, and integration of micro air vehicles (MAVs), heterogeneous MAV swarm has emerged as a powerful platform for cooperative missions in confined and complex environments, such as search-and-rescue operations related to urban disasters, indoor reconnaissance, and small-scale surveillance[1−3]. Due to the severe constraints inherent to MAV platforms, including limited onboard energy, small communication range, low payload capacity, and high mobility, efficient task allocation and coalition formation are critical for mission success. In these scenarios, heterogeneous MAVs exploit their complementary capabilities (e.g., different sensor types, endurance, or payload configurations) to form coalitions, enabling the swarm to accomplish complex, resource-intensive tasks that individual MAVs cannot handle alone. For example, in search-and-rescue operations related to urban disasters, MAVs equipped with thermal cameras, communication relays, and lightweight manipulators can collaboratively cover the search areas, relay data, and deliver small payloads, significantly improving the response time and mission effectiveness under tight resource constraints.

The core challenge for MAV swarms is that of dynamically organizing their agents to meet the task demands while adapting to severe physical and communication constraints. Coalition formation is a key mechanism in this process, directly determining the task allocation efficiency, overall swarm performance, and system resilience in highly dynamic environments. Therefore, developing coalition formation strategies tailored to the unique characteristics of MAVs, such as short range, intermittent communication, and rapid topology changes caused by agile flight maneuvers, holds substantial theoretical importance and practical value for advancing cooperative MAV operations[4].

In recent years, the academic community has conducted extensive research on coalition formation and has achieved remarkable progress especially in disjoint coalition formation (DCF)[5]. This method partitions agents into non-overlapping coalitions, generating stable collaborative structures for multiple distinct tasks. Its advantages include simplifying task allocation processes, reducing resource conflicts, and improving system manageability. For example, Mahdiraji et al.[6] formalized coalition formation as a combinatorial optimization problem and proved it to be NP-hard, laying the theoretical foundation for subsequent studies. More recently, DCF mechanisms have been extensively optimized for specific industrial and defense applications. For instance, Dai et al.[7] proposed a coalition game-theoretic strategy for distributed radar networks, whereby radars are partitioned into disjoint groups to optimize power allocation and enhance multi-target detection performance. In the domain of intelligent transportation, Maxim et al.[8] applied a coalitional distributed model predictive control approach, dynamically organizing vehicles into stable, non-overlapping platoons to maintain the formation stability under switching communication topologies. Furthermore, Wang et al.[9] addressed the coalition structure generation problem in edge computing environments by developing an arbitrary discrete political optimizer to efficiently partition edge nodes for handling concurrent multitasking workloads.

However, the limitations of DCF have become increasingly apparent[10,11]. First, it assumes agents participate in only one coalition, leading to suboptimal resource utilization and failing to fully exploit the multifaceted capabilities of heterogeneous agents, especially in resource-scarce or task-diverse scenarios[12,13]. Furthermore, the mobility of MAVs leads to rapid time-varying communication topologies, resulting in frequent link breaks, increased delays, and inconsistent situational awareness, challenges that most existing DCF methods do not adequately address in real MAV operations.

To address these issues, overlapping coalition formation (OCF) has emerged as a promising approach[14], allowing agents to participate in multiple coalitions simultaneously, thereby fully utilizing their diverse resources and capabilities, improving resource efficiency, and enhancing system flexibility and robustness in dynamic environments[15]. For example, in heterogeneous MAV networks, a single MAV can provide sensor data for reconnaissance tasks while serving as a relay for communication tasks, thus achieving parallel multi-task execution through overlapping coalitions and optimizing system performance. Research has advanced OCF for resource allocation in heterogeneous agent networks, among others. Qi et al.[16] proposed a sequential OCF game that considers the overlapping and complementary relations of resource properties as well as the task execution order, introducing a bilateral mutual benefit transfer order to optimize cooperative task resource allocation. However, this negotiation mechanism inherently relies on the assumption of a stable, globally connected network to reliably facilitate these mutual benefit transfers. Similarly, Li et al.[17] developed a heuristic OCF method based on game theory, integrating task priority, dynamic task benefits, and discrete resource constraints into the resource matching process. Although it is highly effective in idealized settings, their model fundamentally assumes instantaneous, delay-free communication to achieve global information synchronization and global priority sorting. To adapt to increasingly complex operational requirements, recent studies have explored OCF mechanisms. Chen et al.[18] offered a capability-centric analysis of coalition formation, extracting key problem elements using the 5W1H method (objects, coalitions, missions, capabilities, and environments) and providing a multi-view analysis through logic diagrams and structure charts; they constructed a general mathematical model to demonstrate how these concepts contribute to capability aggregation in CF, highlighting their applications in multi-robot and human-robot collaboration. Ai et al.[19] proposed an OCF-enabled noncooperative game combined with multi-agent deep reinforcement learning for UAV-assisted resource allocation. Similarly, Zhou et al.[20] designed a dual-rationality-guided partially OCF game approach to jointly optimize task and spectrum in unmanned swarms.

Despite these developments, the majority of the existing OCF studies still overlook the impact of time-varying topologies on communication range and delay, which is particularly severe for MAVs due to their small size, low transmission power, and agile flight patterns. This limitation significantly restricts the applicability of current methods in realistic MAVs' search-and-execute missions, where swarm members must discover unknown task locations through collaborative search, detect targets within short sensor ranges, and form coalitions on-the-fly under constrained and fluctuating connectivity[21].

To address the abovementioned critical challenges in heterogeneous MAV swarms executing multiple resource-demanding tasks under time-varying communication topologies, this paper proposes a capability-centric OCF method specifically designed for MAV platforms. The approach emphasizes dynamic capability aggregation and alignment between MAVs' heterogeneous onboard resources and task requirements, enabling the formation of robust, adaptive overlapping coalitions that can respond to rapid network changes. It integrates a market-based negotiation framework with capability-driven heuristics and explicitly incorporates communication constraints (range, delay, and topology dynamics) to ensure practical feasibility. The main features of this method are as follows:

(1) It is a novel capability-centric model for OCF tailored to dynamic MAV swarm environments, which facilitates efficient capability aggregation, maximizes resource utilization under tight constraints, and provides robustness against topology variations and inconsistent situational awareness;

(2) The Topology-Adaptive Capability Matching (TACM) algorithm is specifically developed to solve OCF in time-varying communication environments typical of MAV operations. TACM effectively handles interdependent tasks, communication distance and delay constraints, and dynamic coalition adjustments during search-and-execute missions.

The remainder of the paper is organized as follows. The next section presents the system model and problem formulation. The section "Capability-centric overlapping coalition formation method" details the proposed TACM method. The section "Complexity and Communication Overhead analysis" evaluates the performance of this method through simulations. The last section concludes the paper and suggests future directions.

Problem statement and model formulation

Scenario description

In this paper, we consider a dynamic environment in which a heterogeneous MAS executes search-and-execute missions. The swarm consists of MAVs with diverse capabilities, including varying sensor ranges, communication radii, and resource types (e.g., energy or specialized tools). Initially, the MAVs are unaware of task locations and resource requirements, which specify the types and quantities of capabilities needed.

MAVs perform collaborative search using predefined strategies such as grid-based or swarm-inspired patterns to cover the area. A task is detected when an MAV's distance to the task location falls within its detection radius, revealing the task's position and capability demands. This initiates the OCF for the task. The detecting MAV broadcasts details to MAVs within its time-varying communication topology, with which connections exist only if distances are within communication ranges, incorporating time delays. This leads to local communication and potential inconsistencies in situational awareness due to mobility-induced topology changes.

Coalition formation occurs if the aggregate remaining capabilities in the local group exceed the task's needs. The proposed capability-centric method then forms the coalition, allowing overlapping memberships for efficient capability use. Coalition members synchronize to arrive at the task location simultaneously for execution. The process continues iteratively as tasks are discovered. MAVs in multiple coalitions sequence executions based on priorities or proximity, managing resources across tasks. The mission ends when all tasks are completed or resources are exhausted. This setup enhances the adaptability under communication constraints in uncertain environments.

Capability-centric model for OCF

In this subsection, we formalize the OCF problem through a capability-centric lens, treating MAVs' resources as generalized capabilities that can be aggregated to meet task demands. This approach extends traditional resource-constrained models by emphasizing dynamic capability matching, where MAVs' diverse abilities (e.g., sensing, computation, or mobility) are quantified and combined in overlapping structures to maximize the overall system efficiency while adhering to individual capacity limits and inter-task conflicts.

Consider n MAVs, denoted by the set $\mathcal{A} = \{a_1, a_2, \ldots, a_n\}$, and m tasks, denoted by $\mathcal{T} = \{T_1, T_2, \ldots, T_m\}$. Each task $T_i \in \mathcal{T}$ requires a specific set of capabilities, represented by a demand vector $\boldsymbol{D}_i = [d_1^i, d_2^i, \ldots, d_p^i]^\top \in \mathbb{R}_{\geq 0}^p$, where p is the number of capability types (e.g., energy, processing power, or sensor accuracy). Similarly, each MAV $a_j \in \mathcal{A}$ possesses an initial capability endowment vector $\boldsymbol{R}_j = [r_1^j, r_2^j, \ldots, r_p^j]^\top \in \mathbb{R}_{\geq 0}^p$, which diminishes as capabilities are allocated across tasks.

A coalition $\boldsymbol{C}_i \subseteq \mathcal{A}$ is formed for task T_i, where the MAVs in $\boldsymbol{C}_i$ contribute portions of their capabilities to fulfill $\boldsymbol{D}_i$. Unlike disjoint coalitions, overlapping is permitted, allowing MAV a_j to join multiple coalitions (e.g., $\boldsymbol{C}_i$ and $\boldsymbol{C}_l$ for $i \neq l$), provided the total contributions from a_j do not exceed $\boldsymbol{R}_j$. For each coalition $\boldsymbol{C}_i$, define the contribution matrix $\mathbf{W}_i \in \mathbb{R}_{\geq 0}^{n \times p}$, where the j-th row $\mathbf{W}_{ji} = [w_1^{ji}, w_2^{ji}, \ldots, w_p^{ji}]$ represents the capabilities allocated by MAV a_j to task T_i (with $\mathbf{W}_{ji} = \mathbf{0}$ if $a_j \notin \boldsymbol{C}_i$). The remaining resources of the MAV a_j after completing its assigned tasks can be represented as

$$\begin{cases} \boldsymbol{S}_j = [s_1^j, s_2^j, \ldots, s_p^j]^\top \in \mathbb{R}_{\geq 0}^p \\ s_k^j = r_k^j - \sum\limits_{i=1}^{m} w_k^{ji}, k = 1, \ldots, p \end{cases} \tag{1}$$

The capability aggregation for task T_i must satisfy the demand constraint given below, ensuring the collective contributions meet or exceed the required levels.

$$\sum_{j=1}^{n} \mathbf{W}_{ji} \geq \boldsymbol{D}_i \tag{2}$$

For each MAV a_j, the total allocation across all tasks is bounded by its endowment:

$$\sum_{i=1}^{m} \mathbf{W}_{ji} \leq \boldsymbol{R}_j \tag{3}$$

This constraint prevents overcommitment and resolves potential conflicts arising from overlapping memberships.

The utility of coalition $\mathbf{C}_i$ is defined as a function $v(\mathbf{C}_i)$ that quantifies the value of the task being completed, t_i, incorporating factors such as task priority, timeliness, and efficiency of capability usage. In a capability-centric framework, we adopt a super-additive utility form to reward effective aggregation:

The utility of coalition $\mathbf{C}_i$ in handling task T_i is quantified by its performance metric, which balances the value gained from task accomplishment against the associated expenditures:

$$J(\boldsymbol{C}_i) = u_i - \theta(C_i) \tag{4}$$

where, u_i denotes the inherent reward for achieving task T_i (e.g., a priority-based score or mission payoff). Meanwhile, $\theta(C_i) = \sum_{a_j \in C_i} L_{ji}$ aggregates the operational overheads, specifically the summed distances (path lengths) from each coalition member a_j to the task site T_i, reflecting mobility costs influenced by current positions and topology dynamics.

The goal of an OCF problem is to obtain an overlapping coalition structure (OCS)[6]. An OCS comprises the ensemble $\{\mathbf{C}_1, \mathbf{C}_2, \ldots, \mathbf{C}_m\}$, assigning one coalition per task. The aggregate performance of an OCS is the cumulative metric across all coalitions:

$$J(OCS) = \sum_{i=1}^{m} J(C_i) \tag{5}$$

Define Π_A as the space of all viable OCS configurations, incorporating capability bounds and overlap feasibility. The central aim of the OCF challenge is to identify the superior OCS that elevates the total performance to its peak, cast as an optimization task with the following constraints:

$$\begin{cases} OCS^* = \arg\max_{OCS \in \Pi_A} \sum\limits_{i=1}^{m} J(C_i), \\ \text{Subject to:} \\ \sum\limits_{j=1}^{n} w_k^{ji} \geq d_k^i, (k \in \{1, \ldots, p\}, i \in \{1, \ldots, m\}) \\ \sum\limits_{i=1}^{m} w_k^{ji} \leq r_k^j, (k \in \{1, \ldots, p\}, j \in \{1, \ldots, n\}) \end{cases} \tag{6}$$

The initial constraint ensures that the pooled capabilities for each task at least match its demands, underscoring the capability aggregation principle. The subsequent constraint safeguards against MAV overload by capping the total allocations per capability type, averting disputes in multi-coalition scenarios.

To integrate time-varying topologies, coalition assembly is confined to MAVs in mutual communication proximity, defined by range-limited links and delay considerations, promoting decentralized viability. Compared with non-overlapping alternatives, this setup enables fractional capability commitments, increasing the adaptability in constrained settings: for example, enabling an MAV with ample computational resources to aid several analytical tasks while conserving mobility-related capabilities for others. Furthermore, the capability-centric orientation supports responsive adjustments, where quantitative assessments of capabilities guide allocations amid challenges such as sudden task appearances or inconsistent awareness due to network shifts. This not only aligns with practical deployment in search-and-execute operations but also extends the model's applicability to broader domains such as disaster response or logistics coordination, where balancing rewards and costs under uncertainty is paramount.

Dynamic agent-task matching framework in capability-centric coalition formation

To illustrate the interplay between tasks, capabilities, MAVs, and coalitions in the capability-centric coalition formation framework, Fig. 1 presents a conceptual diagram of the capability-centric dynamic matching mapping for heterogeneous MAS. This layered representation, inspired by the capability-centric analysis in Chen et al.[18], captures the essence of how capabilities serve as the central driver for forming overlapping coalitions, enabling efficient task allocation under resource constraints and time-varying topologies.

The diagram is structured in four interconnected layers, flowing from top to bottom: tasks T_1, T_2,..., T_m at the uppermost level, followed by capabilities r_1, r_2,..., r_p, then MAVs a_1, a_2,..., a_n, and finally coalitions C_1, C_2,..., C_m at the base. Arrows indicate the directional relationships and dependencies, highlighting a multi-to-multi mapping that evolves dynamically over time. Starting from the tasks layer, each task T_k specifies a set of required capabilities, represented as demand vectors, which must be aggregated from the available resources. These capabilities, represented in the second layer as triangular nodes, act as the pivotal elements (quantifiable attributes such as energy, sensing range, or computational power) that bridge task needs with MAV endowments.

The third layer consists of MAVs, shown as hexagonal nodes, each possessing a unique capability profile. The mapping from capabilities to MAVs is many-to-many, allowing a single capability type to

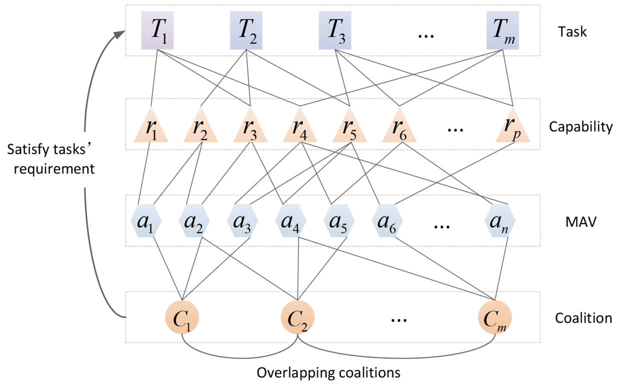

Fig. 1 Conceptual diagram of the capability-centric dynamic matching mapping for MAS.

be provided by multiple MAVs and an MAV to contribute multiple capability types. This flexibility is crucial for overlapping coalitions, as MAVs can allocate partial capabilities to different tasks without full commitment, optimizing the utilization in resource-limited scenarios. The bottom layer represents the coalitions, where groups of MAVs are clustered to collectively meet task demands through capability aggregation. The connections here emphasize that coalitions are not rigid; they can share MAVs, as indicated by the intersecting arrows, to handle multiple tasks concurrently or sequentially.

In this capability-centric view, the mapping process is driven by the aggregation and alignment of capabilities: tasks pull the required capabilities, which in turn select suitable MAVs based on availability and proximity within the current communication topology. The diagram also conveys dynamism: topology variations due to MAV mobility can alter connections, prompting real-time remapping. For instance, if a new task emerges during search, the framework adjusts by reallocating capabilities from nearby MAVs, potentially expanding or merging coalitions to maintain efficiency. Resource constraints are implicitly enforced through the mapping, ensuring no MAV's total contributions exceed its endowment, thus preventing conflicts in overlapping structures.

Capability-centric OCF method

In this section, we introduce the TACM method, a capability-centric approach designed to address OCF in heterogeneous MAS under time-varying topologies. This method integrates a market-based framework with capability-driven heuristics, emphasizing the dynamic aggregation of MAVs' capabilities to meet task demands while accounting for communication constraints. By focusing on capability-centric principles, TACM enables MAVs to partially contribute their resources across multiple coalitions, enhancing system flexibility and robustness in search and execute missions where tasks emerge dynamically.

Overall framework of TACM

The TACM method operates within a distributed architecture, in which heterogeneous MAVs collaboratively form overlapping coalitions through iterative capability matching and negotiation processes. The framework is triggered upon task detection during the search phase: when an MAV's distance to a task falls within its detection radius, it broadcasts the task's position and capability requirements to the nearby MAVs within the current communication topology. This initiates a capability-centric auction process, where MAVs bid based on their available capabilities, considering communication distances and delays to ensure feasible contributions.

The overall workflow of TACM can be summarized as follows:

(1) Task detection and initiation: An MAV detects a task and assumes the role of an auctioneer, evaluating the local group's aggregate remaining capabilities against the task's demand vector. If the capabilities are sufficient, it broadcasts an auction announcement, including the task details and required capabilities.

(2) Capability-based bidding: Eligible MAVs within the communication range submit bids, specifying the portions of their capabilities they can allocate. Bids incorporate capability-centric metrics, such as the efficiency of contribution relative to mobility costs and topology-induced delays.

(3) Winner determination and coalition adjustment: The auctioneer selects winners using a capability aggregation utility function, forming an initial coalition. If the aggregated capabilities fall short, the auction reopens. Overlapping memberships are allowed, with MAVs updating their remaining capabilities post-allocation.

(4) Execution transition: Once a coalition meets the task's capability threshold, members synchronize for simultaneous arrival using established path-planning algorithms. The process iterates for subsequent tasks, with MAVs managing overlapping commitments via priority sequencing.

This framework ensures adaptability to time-varying topologies by limiting interactions to local neighborhoods and dynamically reassessing the availability of capabilities. Compared to traditional disjoint coalition methods, TACM's capability-centric focus promotes efficient resource utilization, as MAVs can fractionally commit capabilities across tasks without overcommitment.

Capability-driven coalition formation algorithm

Task detection and initiation

The initiation and detection of tasks in the capability-oriented coalition formation algorithm play a pivotal role in addressing task appearances within evolving, topology-fluctuating settings involving diverse MAVs. This segment activates an MAV, which pursues its designated exploration path and spots any task that enters its perceptual boundary. Consequently, the spotting MAV appoints itself as the auctioneer for that task, directing the decentralized process of capability consolidation and group assembly.

Owing to the incessant movement of MAVs and the built-in boundaries on signal propagation and postponements, the auctioneer does not possess instantaneous knowledge of its proximate connections or the shifting structural configuration. To regulate the extent of data spread and permit MAVs to assess their feasible participation grounded on capability compatibility, a constrained dissemination threshold, labeled as the hop limit H_{max}, is utilized. Every transmitted packet begins with a hop limit initialized at H_{max}, reducing by one per retransmission. Dissemination terminates when the tally hits zero, thus limiting surplus data flow across the system.

For instance, as illustrated in Fig. 2, consider a scenario where the hop limit is set to $H_{max} = 2$. Upon identifying task T, the auctioneer (MAV A2) initiates a broadcast to its immediate neighbors. In the schematic network layout, intermittent lines denote the functional communication links within the valid signal perimeter. Through this regulated forwarding mechanism, the information is propagated only to MAVs within a specific local region (e.g., a two-hop radius). Conversely, remote entities such as MAV A6 are intentionally excluded, as reaching them requires three hops—exceeding the

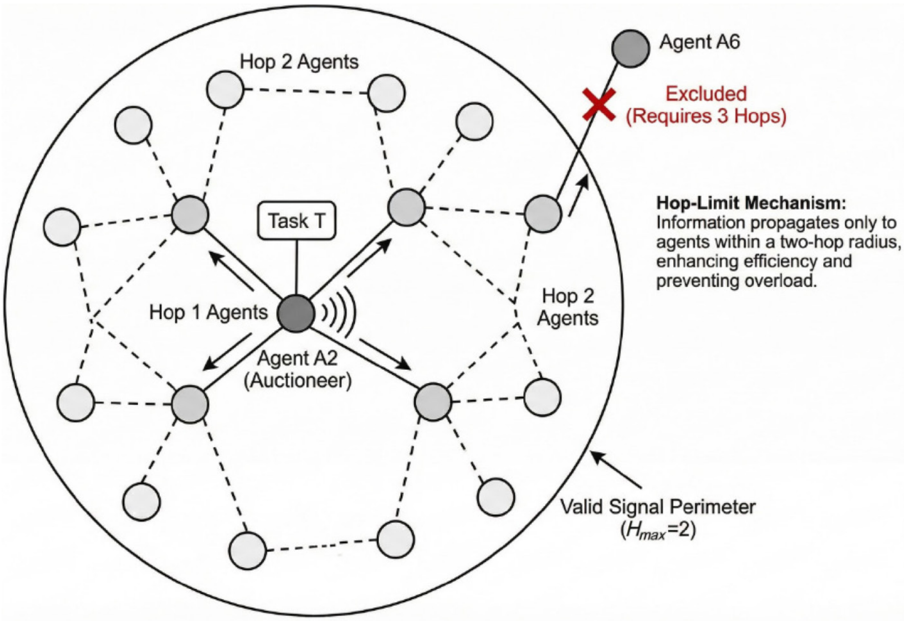

Fig. 2 Schematic diagram of the hop-limited propagation mechanism.

defined threshold. This constraint effectively limits the scope of data dissemination, thereby enhancing the communication efficiency and preventing the network overload.

Alterations in the network during group construction may sever links, risking stagnation in operations. To address this, those receiving signals appraise their dependability as conduits. The schedule for group establishment, covering spread, submission of offers, and choice, sets a maximum duration δ_c for activities, defined by:

$$\delta_c = t_{TA} + 3 \cdot (H_{\max} \cdot \Delta\tau) + \delta_t \tag{7}$$

In this expression, the term $3 \cdot H_{max}$ covers reciprocal packet journeys, t_{TA} denotes the interval for the auctioneer to build the coalition, and δ_t acts as an adjustable cushion for uncertainties.

For an emerging task T_i (with i ranging from 1 to m), detected at instant t_{now}, the anticipated coalition formation moment $t^c_{T_i}$ becomes:

$$t^c_{T_i} = \delta_c + t_{\text{now}} \tag{8}$$

When a certain MAV $a_k, (k \in \{1, .., n\})$ obtains a forwarded notice, it anticipates its placement at $t^c_{T_i}$. If this foreseen location resides inside the initiator's signal span, it considers itself a reliable conduit.

The notice bundle P_j, dispatched by overseer MAV a_j regarding task T_i, is organized as:

$$P_j = [a_j, T_i, f^i_j, t_{\text{now}}, \delta_c, H] \tag{9}$$

where, a_j represents the overseer, T_i transmits the task's placement and capability needs array, f^i_j measures the foreseen capability consolidation advantage for a_j tackling T_i, and H is the leftover forwarding quota (commencing at H_{max}). The advantage indicator f^i_j is computed as follows:

$$f^i_j = V_{ai} \exp(-\beta_i \cdot t_{ETA}(j,i)) \tag{10}$$

where, $t_{ETA}(j,i)$ denotes the journey span for MAV a_j to the task locale, V_{ai} signifies the task's foundational worth, and β_i adjusts the time-based reduction.

When multiple MAVs in intersecting network areas spot tasks at the same time, overlapping notices might emerge, depleting resources and causing discrepancies. To prevent resource depletion and inconsistencies when multiple MAVs in intersecting network areas detect a task simultaneously, a distributed coordination mechanism is triggered. Specifically, this mechanism operates as a distributed max-consensus protocol. Let $\mu^{(t)}_{sync,i}$ denote the highest advantage metric known to MAV a_i at iteration round t. In each round, MAV a_i exchanges its current metric with its one-hop neighbors $\mathcal{N}_i$ and updates its state according to the expression $\mu^{(t+1)}_{sync,i} = \max_{j \in \mathcal{N}_i \cup \{i\}} \{\mu^{(t)}_{sync,j}\}$.

Regarding convergence, because the data dissemination is fundamentally constrained by the hop limit H_{max}, the maximum network diameter of the interacting local topology is bounded by $2H_{max}$. Consequently, the max-consensus algorithm is theoretically guaranteed to reach aglobal agreement within this local subset if the number of synchronization rounds satisfies the convergence condition $N_{sync} \geq 2H_{max}$.

Across these N_{sync} cycles, the local system deterministically aligns with the notice boasting the superior advantage. This process highlights the TACM algorithm's efficiency in resolving conflicts: when multiple simultaneous detections occur, conflicts are resolved strictly within the local sub-network via a finite, small number of scalar comparisons. This rapidly yields a single, uniquely determined auctioneer without requiring global swarm flooding or central oversight, thereby ensuring efficient synchronization and durability amidst structural variations.

Anchored in capability-focused appraisal, this segment promotes selective spread while adapting to ambient changes, laying a solid foundation for following offer and choice phases in crafting flexible, intersecting coalitions.

Capability-based bidding

Following the task detection and initiation phase, the capability-based bidding stage enables MAVs to propose their partial contributions to the coalition, aligning with the capability-centric paradigm of TACM. This phase is activated upon receipt of an auction announcement, in which each eligible MAV evaluates its ability to commit portions of its remaining capabilities, factoring in communication delays, mobility costs, and potential overlaps with other tasks. By incorporating topology-aware adjustments, this bidding mechanism ensures that the proposals reflect realistic feasibility in dynamic environments, promoting efficient capability aggregation across overlapping coalitions.

When an MAV a_k receives the announcement for task T_i, it first verifies the stability of its connectivity and compatibility of its capability. It then forecasts its trajectory to confirm it can maintain a viable link with the auctioneer throughout the anticipated coalition formation duration δ_c. If stable, U_k computes its remaining capabilities $\boldsymbol{S}_k$ that it can allocate to T_i without exceeding endowments or conflicting with prior commitments. Then, the MAV a_k formulates and sends a bid message B_k to the auctioneer:

$$B_k = [a_k, T_i, f_k^i, \boldsymbol{S}_k, t_k^i, t_{ETA}(k,i)] \tag{11}$$

where, a_k is the bidder's identifier, T_i references the task, f_k^i is the benefit, $\boldsymbol{S}_k$ details the remaining capabilities of the MAV a_k, t_k^i is the bid timestamp, and $t_{ETA}(k,i)$ is the arrival estimate. The bid is relayed through stable intermediaries if direct communication is unavailable, with the relay counter ensuring that the propagation stays within H_{max}.

The auctioneer collects bids until the deadline. This bidding phase not only retains robustness against inconsistent situational awareness but also advances capability-centric optimization, enabling MAVs to propose tailored, topology-resilient contributions.

Winner determination and coalition adjustment

In the capability-driven coalition formation algorithm, the winner determination and coalition adjustment phase follows the bidding stage, where the auctioneer evaluates the collected proposals to select optimal contributors and refine the overlapping coalition. This phase ensures that capability aggregation maximizes the task utility while adapting to communication constraints and dynamic topologies. By incorporating a greedy heuristic with topology-sensitive penalties, it efficiently builds coalitions, allowing for partial capability commitments and iterative adjustments to handle incomplete formations or topology shifts.

Upon the expiry of the bidding deadline, adjusted for maximum delays in the current topology, the auctioneer processes all timely bids, discarding any that arrive late to mitigate the inconsistencies arising from latency. For each valid bid from MAV a_k on task T_i, the auctioneer verifies the remaining capabilities $\boldsymbol{S}_k$ against the task's remaining demand $\boldsymbol{D}_i^{remain}$, ensuring no over-allocation.

The core of winner selection employs a capability-centric greedy algorithm to iteratively construct the coalition $\boldsymbol{C}_i$, starting from an empty set with $\boldsymbol{D}_i^{remain} = \boldsymbol{D}_i$. In each iteration, the auctioneer computes the marginal capability value for each bidder:

$$\Delta v_k^i = f_k^i - L_{ki} - \gamma \cdot \eta_{ki} \tag{12}$$

where, f_k^i is the bidder's benefit, L_{ki} is the mobility cost, γ is a delay-weighting factor, and η_{ki} accounts for topology-induced communication latency. This extended metric, newly introduced in TACM, penalizes bids from distant or delayed MAVs, enhancing the robustness in time-varying networks by favoring low latency contributors for better synchronization.

From a theoretical consistency perspective, it is crucial to distinguish this marginal capability value from the overarching coalition utility function defined in Eq. (4). This equation represents the objective operational utility of the executed mission, where costs are strictly bound to the physical mobility and energy expenditure required to reach the task. In contrast, Eq. (12) serves as an online heuristic penalty mechanism used exclusively during the pre-execution negotiation phase. The topology-induced communication latency (η_{ki}) represents the "friction" of distributed decision-making rather than a physical execution cost. By decoupling these two metrics—embedding η_{ki} into the decision heuristic [Eq. (12)] to guarantee robust coalition assembly, while keeping the physical objective function [Eq. (4)] unadulterated—TACM successfully navigates time-varying topologies without distorting the final evaluation of the swarm's operational efficiency.

The algorithm selects the bidder that maximizes the marginal value Δv_k^i, provided that the proposed contribution (remaining capabilities) vector $\boldsymbol{S}_k$ addresses at least one unmet demand type in the task's remaining demand vector $\boldsymbol{D}_i^{remain}$. Upon selection of bidder a_k, the following operations are performed. First, MAV a_k is added to the coalition $\boldsymbol{C}_i$. Second, the remaining demand vector is updated as $\boldsymbol{D}_i^{remain} = \boldsymbol{D}_i^{remain} - \min(\boldsymbol{S}_k, \boldsymbol{D}_i^{remain})$ for each capability type, thereby preventing allocation excess. Third, the contribution $\boldsymbol{S}_k$ is tentatively reserved within MAV's remaining capabilities to ensure commitment feasibility during the adjustment process.

Iterations continue until $\boldsymbol{D}_i^{remain} = \boldsymbol{0}$ or no further positive Δv_k^i bids remain. If the coalition meets the demand, the auctioneer broadcasts authorization messages to the winners:

$$A^j = [T_i, \boldsymbol{C}_i, t^i] \tag{13}$$

where, t^i is the authorization timestamp. Winners confirm receipt, update their capability ledgers to reflect the commitment, and prepare for execution transition.

To address the potential incompleteness—arising from unexpected communication losses with selected bidders, topology changes causing bidder dropouts, or distributed conflicts—an explicit adjustment mechanism is triggered. It is important to acknowledge that as a distributed greedy heuristic, TACM does not mathematically guarantee finding a feasible coalition in a single iteration, even if a valid combination of capabilities exists locally. For instance, in scenarios involving concurrent tasks within overlapping neighborhoods, multiple auctioneers might greedily select the same optimal MAV based on its high marginal value. To respect its endowment constraint, the shared MAV must reject conflicting authorizations, leaving one or more coalitions incomplete.

These inherent limitations of distributed optimization, alongside transient topology-induced delays, strictly necessitate the adjustment mechanism. If the remaining demand $\boldsymbol{D}_i^{remain} > \boldsymbol{0}$ after processing valid bids, the auctioneer reopens the auction by rebroadcasting the updated demand to an expanded neighborhood as new MAVs move into range. This iterative process organically resolves distributed conflicts, ensuring the coalition can dynamically recover from link failures.

However, to prevent infinite loops and resource deadlocks when capability demands fundamentally cannot be met or when conflicts cannot be resolved within a reasonable timeframe, we introduce a strict termination condition. The auctioneer maintains a counter for adjustment cycles. If the required capabilities remain unfulfilled after a predefined maximum number of retries, the coalition formation for task T_i is explicitly terminated. The task is temporarily flagged as unachievable, and any tentatively reserved capabilities within the incomplete coalition are immediately released back to the respective MAVs' available ledgers. This termination protocol ensures that MAVs are promptly freed from infeasible tasks and can reallocate their resources to other emergent opportunities, preserving the overall efficiency against situational inconsistencies while advancing capability-centric goals.

Execution transition

The execution transition phase in the capability-driven coalition formation algorithm commences once the auctioneer confirms that the coalition satisfies the task's capability demands, marking the shift from formation to operational deployment. This phase focuses on synchronizing coalition members for task fulfillment, emphasizing simultaneous arrival at the task location while continuously monitoring for emerging tasks amid ongoing mobility. By integrating capability-centric validation, it ensures that allocated capabilities are effectively utilized, with MAVs updating their remaining endowments post-execution to support ongoing overlaps.

Following authorization confirmation from all winners, the coalition members initiate path planning to achieve concurrent arrival, leveraging established algorithms such as the distributed cooperative particle swarm optimization method[20], which generates paths that meet kinematic constraints, ensure simultaneous arrival, and avoid collisions for multi-UAV formation rendezvous. The auctioneer, or a designated leader within the coalition, computes and distributes paths that comply with the simultaneous arrival constraint, minimizing the total travel time while accounting for individual MAV velocities and obstacles.

During execution, MAVs maintain local capability ledgers, deducting committed resources upon task completion. If a pop-up task emerges during transit, detected by any coalition member, the detecting MAV pauses non-critical execution elements and triggers a new auction for the emergent task, potentially reallocating partial capabilities from the current coalition if priorities warrant. This interruption-driven mechanism ensures system responsiveness, with overlapping MAVs sequencing tasks based on utility scores. This execution transition, bolstered by the interruption protocols, guarantees seamless task completion while preserving capability-centric efficiency.

Calculation of the capability consumption in coalition execution

Upon finalizing the coalition formation for a given task, the participating MAVs collectively distribute the required capabilities among themselves. In typical heterogeneous multi-MAV reconnaissance and execution scenarios, many approaches rely on greedy heuristics to determine individual MAVs' capability usage within the coalition. However, such methods often result in some MAVs exhausting their execution capabilities prematurely, limiting them to observational roles only.

In highly uncertain and evolving environments, MAVs with depleted execution capabilities would not be able to join subsequent coalitions when new tasks arise. This diminishes the overall resilience of the multi-MAV system for combined reconnaissance and execution operations, particularly in situations involving potential MAV losses. Uneven capability usage heightens the vulnerabilities in these contexts.

To mitigate this, we adopt an equitable capability usage strategy, in which MAVs with more number of remaining capabilities contribute more, and those with less number contribute accordingly. Post-task completion, the objective is to equalize the leftover capabilities across MAVs as much as possible. This strategy preserves a baseline level of capabilities for each MAV throughout the operation, better equipping the system to handle operational variability and unpredictability.

For an identified task $T_i \in \mathcal{T}$ demanding d_p^i in capability type p, with the coalition $\boldsymbol{C_i}$ assigned to the task, the subset of MAVs in $\boldsymbol{C_i}$ designated to expend type-p capabilities is defined as set $\boldsymbol{C}_i^p$:

$$\boldsymbol{C}_i^p = \left\{ j \in \boldsymbol{C}_i \mid s_p^j \geq \frac{\sum_{j \in \boldsymbol{C}_i} s_p^j - d_p^i}{n(\boldsymbol{C}_i)} \right\} \tag{14}$$

where, d_p^i indicates the amount of capability type p available for task T_i, s_p^j represents the currently available type p capabilities of MAV a_j, and $n(\boldsymbol{C}_i)$ denotes the count of MAVs in coalition $\mathbf{C}_i$. The capability expenditure w_p^{ji} for MAV a_j relating to capability type p in coalition $\boldsymbol{C}_i$ is formulated as:

$$w_p^{ji} = \begin{cases} 0, & j \notin \boldsymbol{C}_i^p \\ s_p^j - \dfrac{\sum_{j \in \boldsymbol{C}_i^p} s_p^j - d_p^i}{n(\boldsymbol{C}_i^p)}, & j \in \boldsymbol{C}_i^p \end{cases} \tag{15}$$

where, $n(\boldsymbol{C}_i^p)$ is the number of MAVs in the coalition $\boldsymbol{C}_i^p$.

An example scenario is provided here to exemplify the equitable capability expenditure method, wherein the demand for type p capabilities by task T_i is set as $d_p^i = 10$, within the coalition $\boldsymbol{C}_i$ tasked with the operation, comprising 3 MAVs. The amounts of type p capabilities possessed by each MAV in coalition $\boldsymbol{C}_i$ are 9, 7, and 3, respectively. Using Eq. (14), we determine the MAVs that must expend the capabilities as $\boldsymbol{C}_i^p = \{a_1, a_2, a_3\}$. Per Eq. (15), the type p capabilities expended by the three MAVs are 6, 4, and 0, respectively. Thus, following the task execution by the three MAVs, the remaining available type p capabilities stand at 3, 3, and 3, respectively. In contrast, employing a greedy method yields uneven remaining capabilities among the MAVs. This underscores that the equitable expenditure method ensures that MAVs with abundant capabilities expend proportionally more. This expenditure framework, by precisely balancing the usage, supports ongoing coalition overlaps and system durability.

To clarify the theoretical basis for the MAVs' remaining capabilities tending to equilibrium upon multi-task iterative execution, we mathematically model the expenditure dynamics. Let s_p^j denote the present available type p capability of MAV a_j prior to executing task T_i. According to Eq. (15), the remaining capability after the expenditure w_p^{ji}, denoted as $(s_p^j)'$, is calculated as:

$$(s_p^j)' = s_p^j - w_p^{ji} \tag{16}$$

Substituting the expression of w_p^{ji} for any participating MAV $a_j \in \boldsymbol{C}_i^p$ yields:

$$(s_p^j)' = s_p^j - \left(s_p^j - \frac{\sum_{j \in \boldsymbol{C}_i^p} s_p^k - d_p^i}{n(\boldsymbol{C}_i^p)} \right) = \frac{\sum_{j \in \boldsymbol{C}_i^p} s_p^k - d_p^i}{n(\boldsymbol{C}_i^p)} \tag{17}$$

Note that the remaining capability is independent of the specific MAV a_j. It mathematically reduces to a constant mean value for all MAVs in $\boldsymbol{C}_i^p$. Therefore, for any two participating MAVs $a_u, a_v \in \boldsymbol{C}_i^p$, the absolute difference in their updated capabilities immediately becomes zero.

For MAVs with insufficient initial capabilities ($a_j \notin \boldsymbol{C}_i^p$), $w_p^{ji} = 0$, indicating that they are naturally protected from further depletion. Consequently, the overall variance of capabilities within the swarm strictly decreases. Over a sequence of iterative tasks, this mechanism functions as a strict level-balancing protocol, mathematically guaranteeing that the system robustly converges toward a state of capability equilibrium.

Building upon this mathematically proven equilibrium, it is important to analyze how this equitable consumption strategy affects the competitive ratio compared to an optimal offline coalition formation algorithm. An optimal offline approach, based on perfect advance knowledge of all task locations and capability demands, can globally schedule overlaps to strictly minimize the total mobility costs. Our online equitable strategy, operating under dynamic uncertainty, forces MAVs with abundant capabilities to expend proportionally more. Although this may occasionally result in marginally sub-optimal local routing (e.g., selecting a slightly farther MAV to preserve a closer, depleted MAV), it incurs a minor theoretical cost, slightly lowering the upper bound of the competitive ratio in deterministic settings.

However, in unpredictable time-varying environments, greedy online heuristics are highly vulnerable to adversarial task sequences, in which premature exhaustion of key MAVs can lead to complete failure in subsequent coalition formations, drastically plummeting the competitive ratio toward zero. Because our equitable expenditure mathematically guarantees a convergent baseline level of capabilities for the swarm, it sustains ongoing overlaps and system durability. Consequently, this strategy significantly elevates the lower bound of the competitive ratio in worst-case online scenarios, trading marginal short-term optimization for long-term mission resilience.

Complexity and communication overhead analysis

To evaluate the applicability and scalability of the proposed TACM algorithm in large-scale heterogeneous MAV swarms, this subsection systematically analyzes its computational complexity (time and space) and communication overhead, providing a quantitative baseline comparison.

Let N denote the total number of MAVs in the global swarm, p represent the number of capability types required by a task, and N_{local} define the subset of MAVs within the viable signal perimeter constrained by the hop limit H_{max}. In large-scale operations, the hop-limited propagation mechanism ensures that $N_{local} \leq N$.

Computational complexity (time and space)

(1) Time complexity: The auctioneer processes bids from N_{local} candidates. Sorting and evaluating the marginal values of these bids across p capabilities yields a worst-case time complexity of $O(N_{local} \log N_{local} + N_{local} \cdot p)$. Because the computation is strictly bounded by the local neighborhood size N_{local} rather than the global swarm size N, the time complexity remains highly scalable.

(2) Space complexity: The auctioneer only needs to store the capability vectors and state information of the local bidders. Thus, the space complexity is strictly $O(N_{local} \cdot p)$, effectively preventing memory overflow to resource-constrained MAVs.

Communication overhead

Beyond computational metrics, bandwidth consumption is critical for MAV swarms. In traditional pure flooding-based approaches, a task announcement is blindly retransmitted by every node, resulting in a message overhead that scales proportionally with the global network edges, leading to a worst-case overhead of $O(N^2)$ in dense configurations. This causes severe network congestion as the swarm scales up. Conversely, TACM explicitly utilizes the hop limit H_{max} to terminate surplus data flow. The information propagates only within the local region, restricting the message overhead per task detection to be strictly bounded by $O(N_{local}^2)$. Since N_{local} is determined by local spatial density and H_{max}, the communication overhead remains constant and highly scalable regardless of the global swarm size N.

Comparison with baseline algorithms

To contextualize the theoretical efficiency of TACM, we quantitatively benchmark its complexity against two representative baseline methods. A comprehensive empirical evaluation comparing their operational performance is detailed in the section "Performance comparison of different coalition formations".

(1) PTCFA (Polynomial Time Coalition Formation Algorithm): As a conventional heuristic approach, PTCFA is highly effective for disjoint coalition assembly in relatively stable environments. Since it evaluates the capabilities across the entire swarm, its time complexity scales with the global swarm size, expressed as $O(N \log N + N \cdot p)$. Consequently, as the swarm size N increases, its computational overhead becomes relatively larger compared to the localized TACM approach.

(2) MOQGA (Multi-Objective Quantum Genetic Algorithm): MOQGA is a powerful population-based metaheuristic optimizer that excels in exploring near-optimal solutions for complex multi-objective task allocation. Its time complexity depends on the population size P and the maximum number of generations G, and is approximately $O(G \cdot P \cdot N \cdot p)$. Although it is highly suitable for scenarios allowing sufficient computation time, its iterative evaluation across the global population introduces a higher computational overhead that grows with the global swarm scale N.

In summary, although each algorithm is tailored for specific operational paradigms, TACM's topology-adaptive bounding mechanism effectively restricts the computational load to the local neighborhood, making it particularly well suited for scalable, real-time responses in time-varying topologies.

Simulation results

To verify the effectiveness and robustness of the proposed TACM method for capability-centric OCF in heterogeneous multi-MAV systems under time-varying topologies, this section presents three sets of simulation experiments. In the following subsection, the TACM method is applied to an illustrative example scenario to exhibit the complete workflow of a dynamic search-and-execute mission executed by heterogeneous MAVs. In the section "Peformance comparison of different coalition formations," the outcomes of our OCF algorithm are compared with other baseline coalition formation methods. In the section "Impact of communication constraints," probabilistic evaluations through Monte Carlo methods are carried out to investigate the effects of communication constraint (e.g., varying delays and ranges) on the performance of the proposed method. Every computational evaluation is executed on a standard desktop computer equipped with an Intel Core i7-13700 processor at 5.2 GHz, with 32 GB of memory and MATLAB 2022b as the development platform.

Illustrative example

To demonstrate the operational workflow of the proposed TACM method, this subsection presents a small-scale illustrative scenario involving six heterogeneous MAVs executing search-and-execute missions for two pop-up tasks in a 3,500 m × 3,500 m dynamic environment. The MAVs exhibit heterogeneity through variations in their carried resources, encompassing different quantities and types of capabilities. Specifically, the two tasks require capability vectors of $\mathbf{D}_1 = [4, 2]$ and $\mathbf{D}_2 = [4, 3]$, respectively, across two capability types. At the outset, the MAVs lack knowledge of task locations and engage in a predefined collaborative search strategy. Detection occurs when an MAV's distance to a task falls within its 500 m sensor range, prompting the detecting MAV to act as the auctioneer and initiate coalition formation for that task using the TACM algorithm. Communication is constrained to a 1,000 m range, incorporating time-varying topologies due to the MAV's mobility and potential delays, ensuring that the MAVs can only interact with those in proximity.

During the mission, the heterogeneous MAVs collaboratively perform search-and-execute operations, navigating the environment according to a predefined grid-based patrolling strategy to locate unknown tasks while adapting to time-varying topologies induced by their mobility. At t = 26.5 s, as illustrated in Fig. 3, in which the triangles denote MAVs and the stars represent tasks, MAV a_2 detects task T_2 upon entering its 500 m sensor range. However, constrained by communication limitations, the current communication radius of a_2 encompasses only MAV a_1 as a reachable neighbor. The aggregated capabilities within this range total [3,3] across the two types, which falls short of fulfilling the resource requirement, [4,3], of task T_2. Consequently, the TACM algorithm cannot establish a viable overlapping coalition for T_2 at this juncture, as the available capability aggregation fails to meet the minimum requirements, prompting continued search efforts or potential topology-driven adjustments in subsequent iterations.

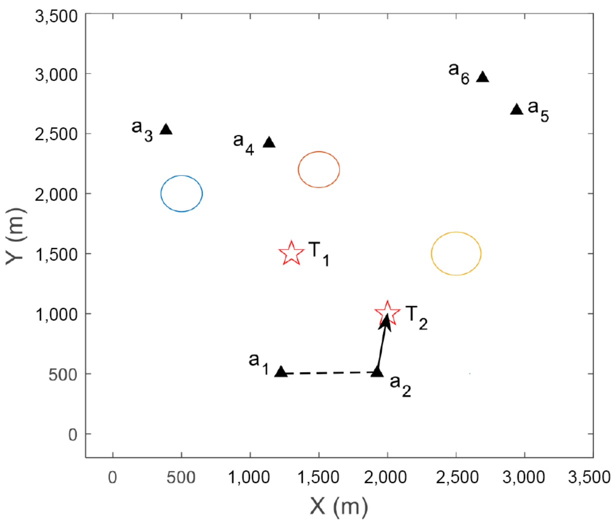

Fig. 3 Configuration of heterogeneous MAVs and tasks at t = 26.5 s in the dynamic environment.

Subsequently, as the mission progresses, at t = 51.1 s, MAV a_1 detects task T_2, thereby assuming the role of auctioneer under the TACM framework. The anticipated communicable MAVs include a_2, a_5 and a_6. The combined capabilities of these MAVs aggregate to [7,5], which sufficiently meets the resource requirement of task T_2, enabling feasible capability aggregation. This triggers the full execution of the TACM algorithm. The resulting coalition task T_2 comprises (a_2,a_5), selected for optimal matching and equitable consumption to preserve residual capabilities for potential future overlaps. Following coalition formation, these MAVs switch to a synchronized path planning for their simultaneous arrival at T_2 by using the distributed cooperative particle swarm optimization method[20] to generate collision-free trajectories compliant with kinematic constraints. Figure 4a illustrates the motion paths of MAVs a_2 and a_5 to T_2; Fig. 4b shows the curvature of their respective paths, ensuring smooth and feasible maneuvers; and Fig. 4c presents the estimated time of arrival, confirming that both MAVs achieve concurrent rendezvous at the task location, thereby validating the method's effectiveness in handling capability constraints and topology variations.

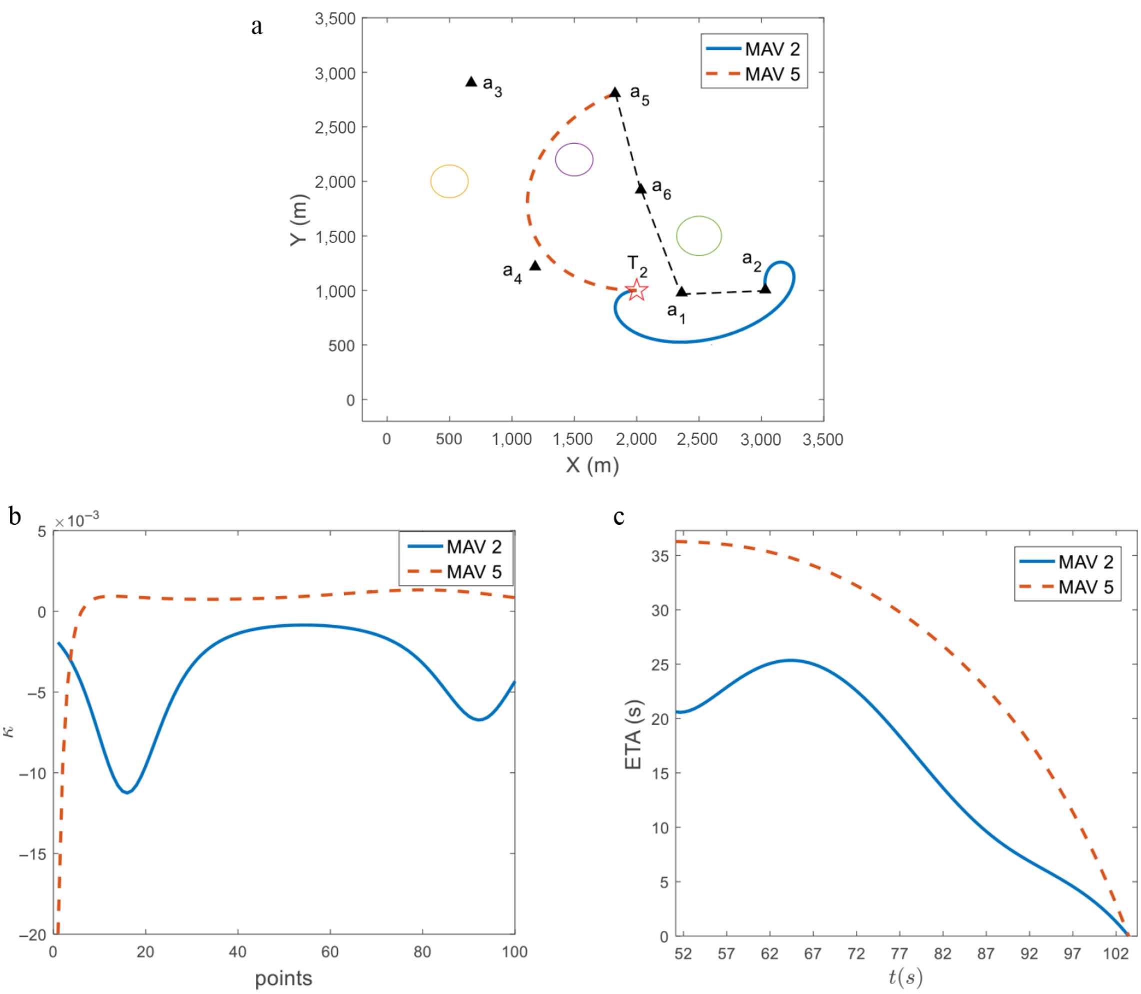

Fig. 4 The coalition formation results for task T_2 at time 51.1 s. (a) MAV coalition for task T_2. (b) Curvature of these paths. (c) Estimated time of arrival.

Later in the simulation, at t = 89.7 s, MAV a_5 identifies task T_1, thereby initiating its role as the auctioneer in the TACM process. Currently, the MAVs accessible for communication are a_1, a_2, a_4, and a_6. The collective capabilities among these MAVs sum to [5,3], adequately surpassing the resource requirement of task T_1. The formed coalition includes (a_1,a_4). After the coalition is established, these MAVs proceed to a coordinated trajectory computation for their synchronized arrival at T_1, as shown in Fig. 5, with the descriptions of the subfigures consistent with those of Fig. 4.

Performance comparison of different coalition formations

To quantitatively assess the performance of the proposed TACM method in capability-centric OCF under time-varying topologies, we conducted a comparative evaluation against two established baselines: PTCFA, which employs a heuristic greedy approach for efficient disjoint coalition assembly, and MOQGA[22], a population-based evolutionary optimizer tailored for multi-objective task allocation in heterogeneous systems. These baselines represent conventional disjoint and metaheuristic strategies, respectively, but lack explicit adaptation to dynamic topologies and overlapping capabilities, making them suitable for benchmarking TACM's innovations. Monte Carlo simulations were employed to quantify the performance in terms of the average mission completion time (total duration from search initiation to the execution of all tasks, including coalition formation and simultaneous arrival) and system utility (aggregated capability fulfillment minus costs, reflecting the overall efficiency and resource leverage).

The simulations are configured within an 5,000 m × 5,000 m area featuring 8 pop-up tasks, each demanding two capability types with random requirements uniformly sampled from 1 to 5 units (inclusive). Heterogeneous MAVs vary in initial capability, starting from a central base and dispersing in diverse directions to execute collaborative search until tasks are detected within their sensor range. Communication constraints are enforced with a 1,000 m radius, and a hop limit H_{max} = 3 is set to simulate realistic topology variations. MAV counts are scaled across 10, 15, 20, 25, and 30 per trial, with each scenario repeated 200 times via Monte Carlo runs to yield statistically robust averages. By keeping the number of tasks constant while varying the swarm size, this setup effectively evaluates the algorithm's performance under different relative task densities (i.e., varying agent-to-task ratios), ranging from resource-scarce environments to resource-abundant ones. Upon detection of a task, coalition formation is triggered per the respective algorithms, followed by path planning for simultaneous arrival using the distributed cooperative particle swarm optimization[23].

The comparative results for average mission completion time and system utility for various MAV numbers are illustrated in Fig. 6a and b, respectively. As shown in Fig. 6a, TACM consistently achieves the lowest mission times across all MAV scales, starting at approximately 298.6 s for 10 MAVs and decreasing sharply to around

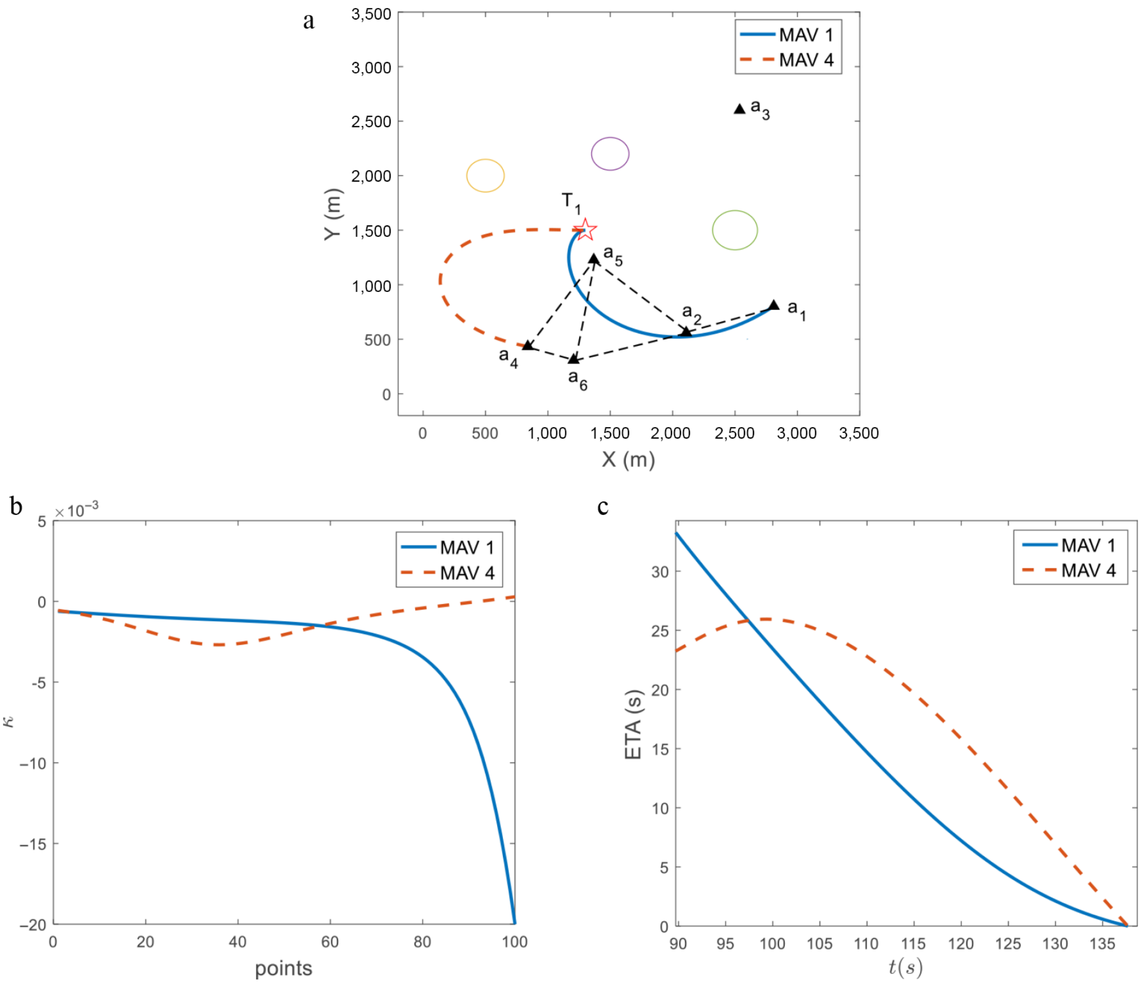

Fig. 5 The coalition formation results for task T_1 at time 89.7 s. (a) MAV coalition for task T_1. (b) Curvature of these paths. (c) Estimated time of arrival.

238.2 s for 30 MAVs, a considerable reduction compared to PTCFA and MOQGA methods. This trend underscores TACM's efficiency in rapid coalition assembly, attributed to its auction-based phases with topology forecasting and marginal capability prioritization, which minimize the delays caused by inconsistent awareness and enable quicker overlaps. In contrast, PTCFA's polynomial-time greedy selection performs adequately for smaller groups but scales less favorably due to its disjoint nature, leading to underutilized resources and frequent re-allocations in dynamic topologies. MOQGA, although robust for multi-objectives, incurs a higher computational overhead because of quantum-inspired iterations, resulting in prolonged formation times exacerbated by limited adaptation to communication constraints.

Figure 6b reveals TACM's dominance in system utility, increasing from 220.2 for 10 MAVs to 248.6 for 30 MAVs, a marginal improvement compared to the PTCFA and MOQGA methods. The progressive gains for TACM stem from its equitable capability consumption and overlapping mechanisms, which maximize resource complementarity and resilience against topology shifts, yielding higher aggregate fulfillment while penalizing fewer delays. PTCFA maintains moderate utility through fast heuristics but suffers from suboptimal overlaps, causing capability wastage in resource-scarce detection. MOQGA's evolutionary search explores diverse structures but converges slower under time-varying networks, limiting its ability to exploit partial cooperations effectively. Overall, these outcomes validate TACM's capability-centric design, demonstrating enhanced scalability, reduced completion times, and superior utility in search-and-execute missions, particularly with increase in MAV density and under time-varying topologies.

Furthermore, the standard deviation error bars in Fig. 6 explicitly indicate the statistical variability across the 200 Monte Carlo runs. The noticeably tighter error margins of the TACM algorithm mathematically demonstrate its superior stability and robustness against the severe randomness induced by time-varying topologies. In contrast, the baseline methods, particularly MOQGA, exhibit higher variance, reflecting their susceptibility to topology-induced deadlocks or sub-optimal allocations.

Impact of communication constraints

To examine how communication delays affect the performance of the proposed TACM method in heterogeneous multi-MAV systems with time-varying topologies, this subsection analyzes systems across different communication delays (0.02 s, 0.1 s, and 1 s) and maximum hop limits (H_{max} ranging from 1 to 5). The evaluations are based on Monte Carlo simulations (200 iterations per setup) in a 5,000 m × 5,000 m mission area containing 8 pop-up tasks, each with two capability types (requirements randomly set between 1 and 5 units). A total of 10 MAVs, heterogeneous in capabilities, initiate from a base and conduct a collaborative grid-based search until tasks are detected within their sensor range. Coalition formation is activated upon detection, with TACM incorporating a hop limit to address inconsistent situation awareness, while maintaining equitable capability distribution for sustained overlaps. Two distinct configurations are tested: (1) a restricted setup with a detection range of 300 m and a communication range of 600 m, emphasizing limited connectivity; (2) an extended setup with a detection range of 500 m and a communication range of 1,000 m, allowing for broader initial interactions.

In the restricted configuration, as presented in Fig. 7a, the system utility displays varied patterns depending on the communication delays. For the minimal delay of 0.02 s, the utility monotonically escalates from $H_{max} = 1$ to $H_{max} = 5$. This is because when the delay is very small, the coalition formation process is short, and the MAVs will not cause significant changes to the communication topology during capability aggregation. In contrast, for the moderate delay of 0.1 s, the utility reaches its peak at $H_{max} = 3$ and then slightly recedes, indicating an optimal balance whereby an additional hop limit enhances bidder diversity but accrues cumulative latency costs. The high delay of 1 s exhibits a more abrupt non-monotonic trend, rising from a low starting point to its maximum at $H_{max} = 2$ before declining sharply at higher hops. Excessive communication delay leads to a longer coalition formation process, making it difficult to move MAVs to form reliable coalitions.

In the extended configuration, as illustrated in Fig. 7b, the trends persist, but with elevated baseline utilities due to larger communicable clusters from the outset. The 0.02 s delay shows a consistent upward trajectory from a low level to a substantially higher level. This is because of the fact that minimal delays result in brief coalition formation durations, preventing significant topology shifts due to MAV movement during aggregation. The 0.1 s delay increases to its peak around $H_{max} = 4$ before dropping, while the 1 s delay maximizes at $H_{max} = 3$ and then falls.

Of particular note in both scenarios is the non-monotonic behavior for delays of 0.1 s and 1 s: the utility initially improves with rising hop limit H_{max}, as hops facilitate wider information dissemination and capability matching, but subsequently diminishes. This occurs

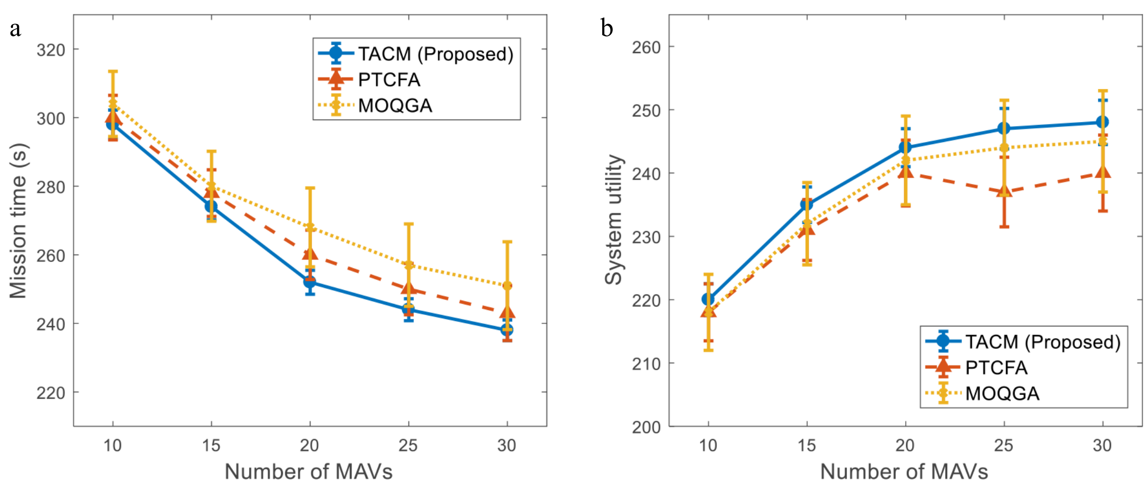

Fig. 6 Comparison of results from three coalition formation methods. (a) Average mission completion time. (b) Average system utility.

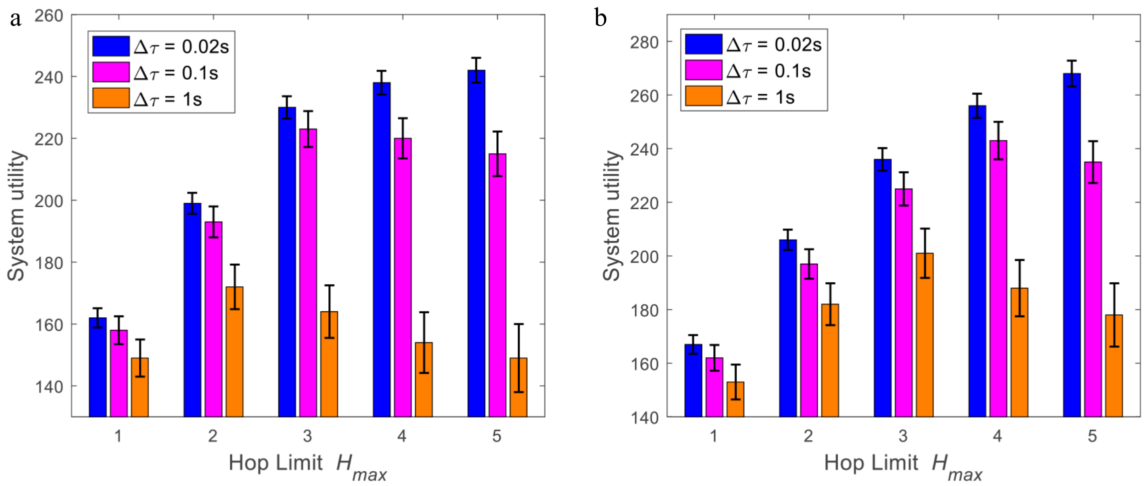

Fig. 7 Impact of communication constraints on the performance of the TACM method. (a) MAV's detection range: 300 m; communication range: 600 m. (b) MAV's detection range: 500 m; communication range: 1,000 m.

because higher delays prolong the coalition formation process, during which MAVs' ongoing movement alters topologies, diminishing the pool of persistently communicable neighbors and potentially causing deadlocks or suboptimal aggregations. In other words, excessive communication delay leads to a longer coalition formation process, making it difficult for moving MAVs to form reliable coalitions. A comparison of Fig. 7a and b reveals that increasing the MAVs' detection range and communication range can effectively counter the increase in communication delay.

The error bars in Fig. 7 represent the standard deviations of the system utility under different network conditions. It is visually evident that while minimal delays yield extremely stable performance with tight variance, higher delays combined with expanded hop limits drastically increase the statistical variability of the results. This statistical evidence confirms that prolonged communication delays severely disrupt the coalition formation process, leading to a highly unpredictable system performance in time-varying environments.

Parameter sensitivity and optimization suggestions

To facilitate the practical deployment of the TACM algorithm in real-world MAV swarms, this subsection provides a sensitivity analysis of the core algorithm parameters—namely, the hop limit (H_{max}) and the delay weighting factor (γ)—along with corresponding optimization suggestions.

Hop limit (H_{max})

As evidenced by the Monte Carlo simulations above, the algorithm's performance is highly sensitive to H_{max}. This parameter governs the trade-off between the scope of capability aggregation and the risk of topology-induced failures.

Optimization suggestion: H_{max} should be dynamically optimized based on real-time network conditions. In environments experiencing severe interference or high transmission delays (e.g., ≥ 0.1 s), H_{max} should be strictly constrained (e.g., H_{max}) to guarantee coalition stability and prevent formation deadlocks. Conversely, in benign environments with minimal latency, H_{max} can be safely increased to maximize the diversity of overlapping candidates and elevate the overall system utility.

Delay weighting factor (γ)

The parameter γ, embedded within the marginal capability value function (Eq. [12]), serves as a heuristic penalty weight during the online coalition formation phase. It modulates the algorithm's sensitivity to topology-induced communication latency relative to physical mobility costs.

Optimization suggestion: The selection of γ should align with the overarching mission priorities. For time-critical missions requiring tight multi-MAV synchronization, or when operating in highly turbulent environments where communication links frequently break, a larger γ is recommended. This setup heavily penalizes high-latency bidders, favoring MAVs with robust, low-hop connections to guarantee successful coalition assembly. Alternatively, if the mission is strictly endurance-bound (where minimizing the physical flight distance is paramount to maximize the final physical utility) and the swarm topology is relatively stable, a smaller γ is optimal.

Conclusions

This paper proposes TACM, a capability-centric OCF method, for heterogeneous MAS under time-varying topologies. By integrating a market-based framework with a hop-limited propagation mechanism and a marginal-value-based algorithm, TACM efficiently forms coalitions while adapting to severe communication constraints. Additionally, an equitable capability consumption strategy is proposed to balance resource depletion in order to enhance the swarm's long-term operational resilience.

The simulations demonstrate that TACM outperforms established baselines (PTCFA and MOQGA), consistently achieving shorter mission completion times and higher system utility. The results confirm that TACM effectively mitigates the performance degradation resulting from transmission delays and topology fluctuations, highlighting the importance of dynamic capability aggregation and topology awareness in robust search-and-execute missions.

Future research will address two main limitations. First, to move beyond the current first-discovered approach, we plan to integrate a dynamic priority-ranking mechanism for preemptive resource reallocation toward emergent high-priority tasks. Second, we will transition to hardware-in-the-loop simulations and small-scale physical flight tests to verify TACM's robustness under real-world aerodynamic and hardware-induced communication constraints.

Author contributions

The authors confirm their contributions to the paper as follows: study conception and design: Yan F, Liao X; data collection: Liao X, Yang T; analysis and interpretation of the results: Liao X, Fu W, Cheng G; draft manuscript preparation: Liao X, Yang T. All authors reviewed the results and approved the final version of the manuscript.

Data availability

The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request.

Acknowledgments

This research was funded by the Scientific Research Program funded by the Education Department of Shaanxi Provincial Government (Program No. 24JK0664).

Conflict of interest

The authors declare that they have no conflict of interest.

Dates

Received 2 January 2026; Revised 12 March 2026; Accepted 14 April 2026; Published online 27 June 2026

References

[1] Ren Y, Zhu F, Lu G, Cai Y, Kong F, et al. 2025. Safety-assured high-speed navigation for MAVs. *Science Robotics* 10(98):6187

[2] Fang C, Mao J, Li D, Wang N. 2024. A coordinated scheduling approach for task assignment and multi-agent path planning. *Journal of King Saud University - Computer and Information Sciences* 36(1):101930

[3] Thelasingha N, Julius AA, Humann J, Humann J, Reddinger JP, et al. 2024. Iterative planning for multi-agent systems: an application in energy-aware UAV-UGV cooperative task site assignments. *IEEE Transactions on Automation Science and Engineering* 22:3685−3703

[4] Huo X, Zhang H, Wang Z, Huang C, Yan H. 2024. A self-learning approach to heterogeneous multi-robot coalition formation under uncertainty. *IEEE Transactions on Automation Science and Engineering* 22:3445−3457

[5] Mali G, Misra S. 2023. CORA: cooperative communication and optimal resource allocation in multihop wireless multimedia sensor networks. *IEEE Internet of Things Journal* 11(3):4076−4084

[6] Mahdiraji HA, Razghandi E, Hatami-Marbini A. 2021. Overlapping coalition formation in game theory: a state-of-the-art review. *Expert Systems with Applications* 174:114752

[7] Dai X, Shi C, Wang Z, Zhou J. 2023. Coalition game theoretic power allocation strategy for multi-target detection in distributed radar networks. *Remote Sensing* 15(15):3804

[8] Maxim A, Pauca O, Caruntu CF. 2024. Coalitional distributed model predictive control strategy with switching topologies for multi-agent systems. *Electronics* 13(4):792

[9] Wang X, Dang J, Zhao S, Guo D, Wang Y. 2022. Coalition structure generation in edge computing environment with multitasking concurrency. *IEEE Internet of Things Journal* 10(5):4324−4338

[10] Sarkar S, Curado Malta M, Dutta A. 2022. A survey on applications of coalition formation in multi-agent systems. *Concurrency and Computation: Practice and Experience* 34(11):e6876

[11] Zhang T, Wang Y, Ma Z, Kong L. 2023. Task assignment in UAV-enabled front jammer swarm: a coalition formation game approach. *IEEE Transactions on Aerospace and Electronic Systems* 59(6):9562−9575

[12] Ng JS, Lim WYB, Dai HN, Xiong Z, Huang J, et al. 2020. Joint auction-coalition formation framework for communication-efficient federated learning in UAV-enabled Internet of Vehicles. *IEEE Transactions on Intelligent Transportation Systems* 22(4):2326−2344

[13] Wei Z, Li B, Zhang R, Cheng X, Yang L. 2023. TBOMC: a task-block-based overlapping matching-coalition scheme for task offloading in vehicular fog computing. *IEEE Internet of Things Journal* 10(17):15209−15222

[14] Ruan L, Li G, Dai W, Tian S, Fan G, et al. 2021. Cooperative relative localization for UAV swarm in GNSS-denied environment: a coalition formation game approach. *IEEE Internet of Things Journal* 9(13):11560−11577

[15] Chen W, Zhao S, Zhang R, Yang L. 2020. Generalized user grouping in NOMA based on overlapping coalition formation game. *IEEE Journal on Selected Areas in Communications* 39(4):969−981

[16] Qi N, Huang Z, Zhou F, Shi Q, Wu Q, et al. 2023. A task-driven sequential overlapping coalition formation game for resource allocation in heterogeneous UAV networks. *IEEE Transactions on Mobile Computing* 22(8):4439−4455

[17] Li Y, Zhang Z, He Z, Sun Q. 2024. A heuristic task allocation method based on overlapping coalition formation game for heterogeneous UAVs. *IEEE Internet of Things Journal* 11(17):28945−28959

[18] Chen J, Guo M, Xin B, Wang Q, Lu S, et al. 2024. Coalition formation problem: a capability-centric analysis and general mode. *Science China Information Sciences* 67(11):212202

[19] Ai B, Ye G, Wu Z, Sun Y. 2025. Overlapping coalition formation-enabled noncooperative game-combined multi-agent DRL for UAV-assisted resource allocation. *IEEE Transactions on Sustainable Computing* 10:29−41

[20] Zhou H, Jia L, Chu F, Qi N, Zhang L. 2026. Joint optimization of task and spectrum in autonomous swarm: a dual-rationality-guided partially overlapping coalition formation game approach. *IEEE Internet of Things Journal* 13:14160−14174

[21] Wang D, Liu J, Lian J, Dong X, Wang W. 2024. Distributed Nash equilibrium seeking for multiple coalition games by coalition estimate strategies. *IEEE Transactions on Automatic Control* 69(9):6381−6388

[22] Mousavi S, Afghah F, Ashdown JD, Turck K. 2019. Use of a quantum genetic algorithm for coalition formation in large-scale UAV networks. *Ad Hoc Networks* 87:26−36

[23] Shao Z, Yan F, Zhou Z, Zhu X. 2019. Path planning for multi-UAV formation rendezvous based on distributed cooperative particle swarm optimization. *Applied Sciences* 9(13):2621