Wiebke Frenkel and Prof. Dr.-Ing. Bernd-Christian Renner. 2025. Harnessing Stubborn AUVs for Decentralized Decision Strategies in Communication-Limited Environments. The 22nd European Conference on Multi-Agent Systems (EUMAS), Bucharest, Romania.

Harnessing Stubborn AUVs for Decentralized Decision Strategies in Communication-Limited Environments

Wiebke Frenkel[1,2] [0000−0001−8030−6963] and
Bernd-Christian Renner[1,3] [0000−0002−6936−6444]

[1] Hamburg University of Technology
Institute for Autonomous Cyber-Physical Systems,
Harburger Schloßstr. 28, 21079 Hamburg, Germany
https://www.tuhh.de/acps/
[2] wiebke.frenkel@tuhh.de
[3] christian.renner@tuhh.de

Abstract. Swarms of autonomous underwater vehicles (AUVs) offer a promising solution for complex missions in dynamic environments with limited communication. Unforeseen events, such as currents, equipment failures, or energy shortages, can prevent some AUVs from completing tasks, requiring rapid redistribution among the remaining vehicles. Decentralized coordination under such conditions is challenging, especially without reliable feedback on overall mission progress. This work investigates a decentralized decision-making process for task reallocation, grounded in opinion dynamics. A key parameter is an agent's *stubbornness*, capturing the tendency to persist with an opinion despite uncertainty. Stubbornness influences when and how the swarm reassigns tasks. Simulations across cooperation strategies and swarm compositions show that allowing controlled divergence in agent opinions, rather than enforcing rapid consensus, improves mission success and resilience. Diverse viewpoints reduce groupthink, encourage exploration, and support flexible adaptation to environmental changes. These results highlight the potential of simple, decentralized heuristics like adaptive stubbornness for robust coordination in uncertain underwater environments.

Keywords: AUV Swarm · Opinion Dynamics · Self Organization · Swarm Intelligence

1 Introduction

Autonomous underwater vehicles (AUVs) are increasingly deployed in collaborative missions, such as environmental monitoring [1] and object tracking [28]. These missions require close coordination among multiple AUVs, which face communication constraints, lack central control, and operate in dynamic environments with changing currents, moving targets, and evolving terrain. A key challenge is decentralized task allocation and reassignment in response to local observations and unforeseen events, requiring agents to make autonomous decisions under incomplete information.

We propose a decentralized strategy for dynamic task redistribution and local mission assessment in communication-limited AUV swarms, grounded in distributed opinion dynamics. A central element is *stubbornness*, a tunable parameter that governs how strongly an agent adheres to its current tasks versus adapting to peer input. Highly stubborn agents maintain their decisions, while less stubborn agents are more responsive to local feedback. Stubbornness enables agents to autonomously balance task commitment and collective flexibility, even under low-bandwidth or intermittent communication.

From this, three research questions arise:

- How can decentralized agents infer whether the collective mission is succeeding or requires adaptation despite limited communication and no global oversight?
- How can stubbornness balance individual task commitment and collective adaptability in decentralized AUV swarms?
- How does the dynamic adjustment of stubbornness based on local observations affect task allocation and mission success?

This work advances the understanding of how simple, decentralized heuristics, such as adaptive stubbornness, support robust coordination in AUV swarms under uncertainty and dynamic conditions, connecting theoretical models of collective decision-making with practical underwater missions.

2 State of the Art

Decentralized coordination of AUV swarms in communication-constrained environments is an increasingly critical topic in swarm robotics research [3]. Unlike unmanned aerial vehicles (UAVs), AUVs rely on acoustic communication, which suffers from low bandwidth, high latency, and frequent dropouts [12], making centralized planning or continuous data exchange infeasible.

To address these constraints, recent work has explored decentralized task allocation methods that can tolerate communication loss. For instance, [24] presents a UAV coordination framework based on onboard visual data and localized decision-making. However, such methods often rely on rich sensory input, particularly visual feedback, which is typically unavailable or unreliable in underwater scenarios due to turbidity and limited visibility. This limitation reduces their applicability to AUV systems and motivates the need for strategies that rely on minimal communication.

Machine learning techniques, particularly deep reinforcement learning (DRL) [6,10] or multi-agent reinforcement learning (MARL) [14], have shown promise in enabling adaptive AUV behavior in complex, uncertain environments. While effective in simulation and controlled environments, DRL-based methods require large amounts of training data, are often sensitive to environment-specific features, and demand significant onboard computational resources. These limitations hinder their effectiveness for making real-time decisions during resource-constrained underwater missions.

Several decentralized coordination strategies have been proposed alongside learning-based approaches. In [15], they introduced a method that combines market-based task allocation with hierarchical task network (HTN) planning to coordinate complex missions with temporal and causal dependencies. While effective, this method requires multiple communication rounds per task and a complete HTN specification, which can lead to delays and increased overhead in dynamic, communication-limited underwater environments. To address these issues, we suggest a lightweight, locally decision-driven strategy that minimizes communication needs and avoids complete prior task modeling.

Opinion dynamic models offer a compelling alternative for enabling agents to reach collective decisions through local interactions without global oversight. Researchers find these models to be an efficient substitute for data-intensive methods in decentralized decision-making. Recent studies [18,23] link sociological opinion formation models to swarm coordination, emphasizing stubbornness as a key factor for filtering out faulty signals and maintaining stability. Agents adjust their beliefs based on information from neighbors, and research [26,27] shows that stubborn agents can prevent premature consensus while preserving informational diversity in uncertain conditions.

Building on these insights, we develop a behavior-based coordination strategy that integrates adaptive stubbornness into the self-organization process of AUV swarms. We model the swarm as a dynamic opinion network where local interactions are leveraged to estimate global mission progress and adapt task allocation accordingly.

3 Problem Statement

We expand on our earlier architecture [9], which includes a central base station for offline mission planning and an AUV swarm. The AUVs navigate to specific predefined sampling waypoints in this simplified underwater mission. To emphasize decentralized coordination, we will not consider sampling energy and will assume that propulsion is the primary energy consumer.

3.1 Communication

In AUV swarms, communication extends beyond simple data exchange, playing a crucial role in the decision-making process. Due to limited bandwidth and intermittent connectivity, agents must determine locally whether and when their information is significant enough to share. This means that communication becomes a strategic action, influencing decisions such as whether to continue or abort a mission, redistribute tasks, or prioritize support for individual agents.

In our work, we rely on a robust, decentralized time division multiple access (TDMA) protocol [16,7] that exchanges minimal but relevant information, which is a crucial aspect given frequent interruptions. The TDMA assigns each of the agent $n \in \{0, 1, \dots, N-1\}$ a time slot of 30 s per round. A transmission is successful only if the receiver is within range; otherwise, the message is

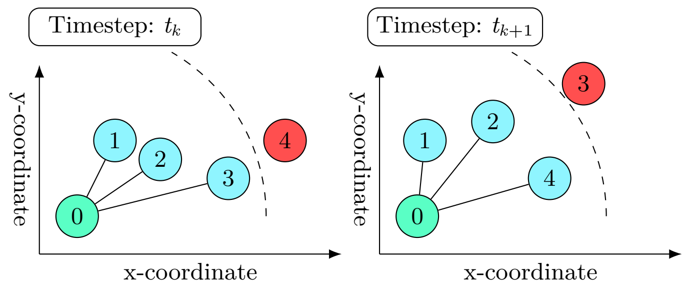

Fig. 1. Local neighborhood of agent 0 (green) at two consecutive timesteps t_k and t_{k+1}. Solid lines indicate direct communication links within the dashed communication range. The changing set of neighbors illustrates the dynamic topology caused by agent movement.

lost. For the sake of simplicity, we do not consider physical transmission effects (such as salinity) and instead focus on inter-agent distance as the primary factor determining connectivity [12].

In fig. 1, we illustrate the dynamic network: agents gain or lose neighbors as the swarm moves, altering the communication graph. This variability poses challenges for robust coordination.

3.2 Energy Considerations in a Greedy Task Redistribution

Energy, primarily consumed by propulsion, limits each AUV's range and endurance in underwater missions. When an AUV lacks sufficient energy to complete its tasks, it must redistribute its workload to nearby agents. To support decentralized task redistribution, each AUV estimates propulsion energy for potential route adjustments using a physics-based model [8,9], calculating the energy cost $E = \mathbf{F} \cdot |\mathbf{x}|$ for a displacement $\mathbf{x}$ under hydrodynamic force $\mathbf{F}$.

We model battery depletion as an equivalent circuit for Li-ion cells [21], where the state of charge (SOC) evolves as

$$\frac{\mathrm{d}}{\mathrm{d}t}\sigma(t) = -\frac{1}{C_{\mathrm{bat}}} \cdot i(t) \tag{1}$$

with the current $i(t)$ derived from the required mechanical power and system losses [8].

The decentralized greedy task redistribution mechanism activates when an AUV can no longer perform its assigned waypoints, triggering reassignment of remaining tasks. Each agent autonomously evaluates redistribution options using local energy estimates and limited peer information.

The assignment of tasks to suitable candidates from the neighbourhood $j \in n; j \neq i$ is represented by directed graphs $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, which are weighted according to segment-wise energy costs. We iteratively test waypoint insertions from the source AUV $i \in n$ into the recipient's route, selecting the position

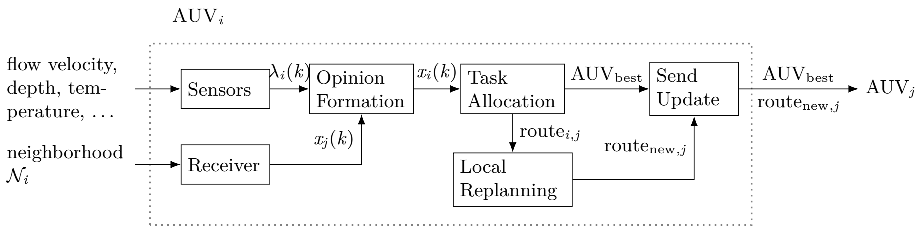

Fig. 2. Schematic representation of the decision-making process within an AUV i.

that minimizes additional energy expenditure. Let $\text{route}_{i\to j} = (v_0, v_1, \ldots, v_k, v_0)$ denote a candidate cyclic path of waypoints from agent i to j. The optimal insertion is then chosen as

$$\text{route}^{\star}_{i\to j} = \arg \min_{\text{route}_{i\to j}} \left(\sum_{\ell=0}^{k-1} E_{v_\ell, v_{\ell+1}} + E_{v_k, v_0} \right), \tag{2}$$

with the sum of segment-wise energy costs along the route.

This greedy, locally informed method supports scalable, resilient coordination under real-world underwater constraints, avoiding the need for computationally expensive global optimizers.

4 The Social AUV Swarm

In our approach, the AUV swarm operates as a dynamic social network [25], where autonomous agents exchange minimal information to collectively manage task allocation and energy constraints. Each AUV continuously evaluates its task load and energy status to decide autonomously when to request or offer help, enabling robust, decentralized coordination under communication and environmental constraints.

In fig. 2 we illustrate the decision-making architecture of a single AUV. Sensor inputs such as flow velocity, depth, and temperature, along with messages received from neighboring AUV $j \in \mathcal{N}_i$, feed into the opinion formation module. This module combines internal assessments with social influence to generate a current opinion on the feasibility of the mission. Based on this, the task allocation module manages workload distribution and local route replanning. Updated task assignments are then communicated to selected peers to enable decentralized swarm coordination.

4.1 Opinion Formation with Stubborn AUVs

In underwater swarm operations, effective decision-making hinges on each AUV's ability to autonomously evaluate its operational state and coordinate with neighbors despite limited communication and incomplete information. To capture this nuanced interplay between individual assessment and social influence, we employ the Friedkin-Johnsen (FJ) opinion dynamics model [23]. Unlike consensus-driven approaches, the FJ model preserves each agent's inherent biases, modeled here as *stubbornness*, enabling resilient, adaptive judgments essential for uncertain and dynamic underwater environments.

Classical models such as DeGroot's [5] drive agents toward complete consensus, which may lead to suboptimal or misleading collective decisions when individual conditions vary widely. By contrast, the FJ model incorporates a stubbornness parameter λ_i that balances an agent's internal judgment with social influence. This mechanism allows agents to consider their disagreement with the neighbors' opinions, thereby enhancing robustness against unreliable or incomplete information [22].

We define the stubbornness of an AUV as a function of its SOC as a measurable quantity for assessing the agent's state to reflect the agent's operational constraints:

$$\lambda_i(k) = 1 - \sigma_i(k), \quad \sigma_i(k) \in [0, 1] \ . \tag{3}$$

Agents with low energy are often convinced that they cannot complete their tasks, which makes them less likely to accept outside suggestions. In contrast, agents with higher energy are generally more open to adapting. This characteristic makes the SOC a straightforward and interpretable indicator of stubbornness in critical decision-making situations.

Each AUV maintains an initial opinion indicating its local evaluation of mission success

$$\tilde{x}_i(k) = 1 - (\xi_0 \cdot \sigma_i(k) + \xi_1 \cdot \mu_i(k)) \ , \tag{4}$$

where $\mu_i(t) \in [0, 1]$ denotes the task completion ratio, and ξ_0, ξ_1 are weighting coefficients. A value of $\tilde{x}_i(k) = 1$ suggests aborting the mission, while $\tilde{x}_i(k) = 0$ favors task continuation.

Following the FJ update rule, each AUV then revises its current opinion by blending its initial assessment with the influence of neighboring agents

$$x_i^{\mathrm{FJ}}(k+1) = \lambda_i(k) \cdot \tilde{x}_i(k) + (1 - \lambda_i(k)) \sum_{j \in \mathcal{N}_i} \omega_{ij} x_j(k) \ , \tag{5}$$

where ω_{ij} are normalized influence weights representing how strongly agent i is affected by the opinion of neighbor j, with $\sum_{j \in \mathcal{N}_i} \omega_{ij} = 1$. This formulation ensures that agents with higher stubbornness maintain a stronger bias toward their evaluation, while more flexible agents consider peer input more.

To mitigate instability due to outdated or missing data, we introduce a healing mechanism

$$x_i(k+1) = \tilde{x}_i(k) + \left(x_i^{\mathrm{FJ}}(k) - \tilde{x}_i(k)\right) \cdot e^{(-\tau \cdot \Delta t)} \ , \tag{6}$$

with τ as a decay rate and Δt as the time since the last successful communication that gradually reverts an agent's opinion toward its initial state when isolated. This dampens the impact of one-sided information and reduces oscillatory behavior in the decision-making process.

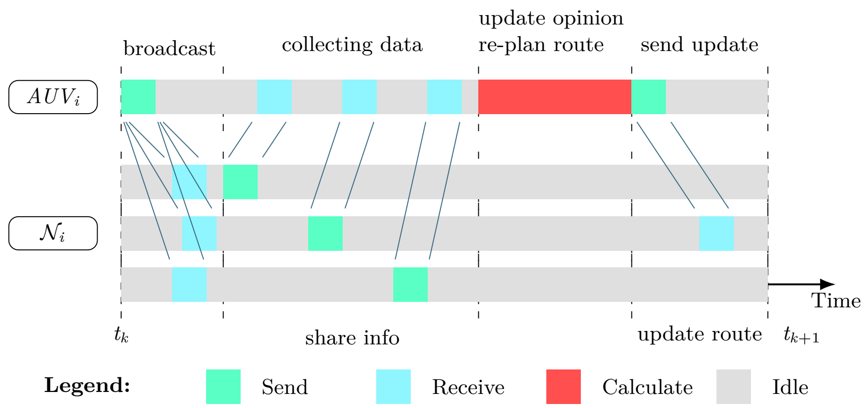

Fig. 3. Sequence of communication and processing steps within a time slot, illustrating the basic cycle of sending, receiving, and recalculating before the next transmission.

4.2 Task Allocation

Unpredictable conditions such as currents, sensor drift, and energy depletion can disrupt task execution in underwater missions. We use a decentralized redistribution mechanism based on peer-to-peer negotiation to maintain robustness.

Each AUV monitors its status and communicates with neighbors. Tasks are reallocated either through help offers from idle agents or help requests from struggling ones. This enables flexible workload balancing without centralized control.

Interaction To ensure collision-free message exchange, we utilize a TDMA protocol section 3.1, where each AUV is assigned a fixed time slot. In fig. 3, we illustrate the processing steps within a single time slot. At the beginning of its slot, the AUV i broadcasts a status message, which may include a help request, an offer, or a regular update, depending on its internal opinion. Neighbors $\mathcal{N}_i$ reply with energy, route, and task data during their slots. Based on this, the AUV updates its opinion and decides whether task redistribution is needed. If it is triggered, the agent i selects a suitable neighbor j and runs a local greedy optimizer section 3.2 to determine which waypoint to reassign and where to insert it in the recipient's route. The selected neighbor j receives the updated route and adjusts accordingly.

A key advantage of our approach compared to auction-based methods is that each sending AUV only needs to query the status of its neighbors once per cycle, rather than repeatedly requesting updated states. This significantly reduces communication overhead and latency. The sending AUV collects all relevant information from neighbors during their assigned time slots and processes it locally to make task redistribution decisions.

Offer If an agent i completes its assigned tasks early, it broadcasts a readiness signal to its neighbors $\mathcal{N}_i$. Each neighbor j is only considered if it still has unfinished tasks. Among these candidates, agent i selects the neighbor j^* with the highest opinion value

$$j^* = \arg\max_{j \in \mathcal{N}_i} x_j(k) \ , \tag{7}$$

where $x_j(k)$ reflects the neighbor's belief that it may not complete its tasks. *Note that a higher opinion value indicates greater pessimism about mission success, i.e., the agent is more likely to request support.* This ensures that help is directed to the agent with the greatest need.

Request When agent i reaches a critical state (e.g., low energy or the willingness to return), it broadcasts a help request. Neighbors evaluate their suitability based on their opinion values $x_j(k)$, and the agent j^* with the lowest opinion (most confident in completing tasks) is selected

$$j^* = \arg\min_{j \in \mathcal{N}_i} x_j(k) \ . \tag{8}$$

Note, a lower opinion value indicates that the agent still believes it can complete tasks, making it a suitable candidate to take over additional work.

5 Evaluation

We analyze how various decentralized behaviors – `Offer`, `Request`, and their combination (`Offer + Request`) – impact AUV swarm performance under communication constraints. These situations necessitate adaptive task redistribution, and we compare them to a non-cooperative baseline.

Each mission starts and ends at the base (0|0). An AUV must return if it reaches a critical state

$$\begin{aligned} \sigma_i(k) < 15\% \ &: \ \texttt{Battery empty} \\ x_i(k) > 0.8 \ &: \ \texttt{Return} \end{aligned} \tag{9}$$

The chosen thresholds ensure safe return and may be dynamically adjusted in future work. A mission is successful if all waypoints are visited and all AUVs return to base.

We run 50 simulated missions for each setup to enable a statistically robust evaluation and account for variability in agent behavior and random waypoint distribution. Our simulations focus on unexpected events causing faster energy drops. To emulate this, some AUVs have a significantly increased SOC variability ($\sigma_i(k) = 20\%$), representing higher mission stress. We evaluate:

- **Minor disturbance**: 30% of AUVs have low SOC. Tests swarm compensation when most agents remain optimistic about continuing the mission.
- **Major disturbance**: Up to 60% of AUVs have low SOC. Assesses resilience when capable agents are in the minority.

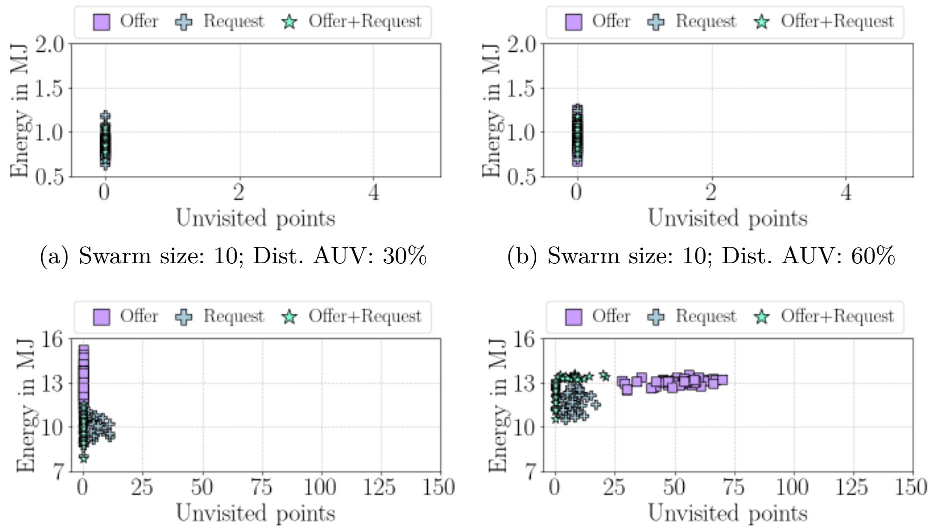

(a) Swarm size: 10; Dist. AUV: 30%
(b) Swarm size: 10; Dist. AUV: 60%
(c) Swarm size: 30; Dist. AUV: 30%
(d) Swarm size: 30; Dist. AUV: 50%

Fig. 4. Propulsion energy in relation to unvisited task waypoints, considering only task-related travel and excluding communication energy.

– **Swarm size variation**: Comparison of a small swarm (10 AUVs, 100 waypoints in a 200×200m area) where most agents can communicate with each other, and a large swarm (30 AUVs, 300 waypoints in a 600×600m area) where the likelihood of inter-agent communication is low.

5.1 Mission Success

This evaluation examines the impact of cooperative task redistribution on mission success, energy efficiency, and execution time in communication-constrained AUV swarms.

Figure 4 illustrates the propulsion energy consumption relative to unvisited task waypoints across different swarm sizes and disturbance levels. In small swarms, all cooperative strategies reliably complete their missions. Notably, they exhibit similar energy consumption levels, indicating that cooperation does not incur significant additional energy costs. For larger swarms under minor disturbance, both the `Offer` and the combined `Offer+Request` strategies complete all missions, whereas the `Request` strategy leaves some waypoints unvisited. Under major disturbance, the limitations of the strategies become evident. The combined `Offer+Request` strategy still completes a substantial portion of the missions, the `Request` strategy fails to cover a significant number of waypoints, and the `Offer` strategy is unable to complete any missions under these highly disturbed conditions.

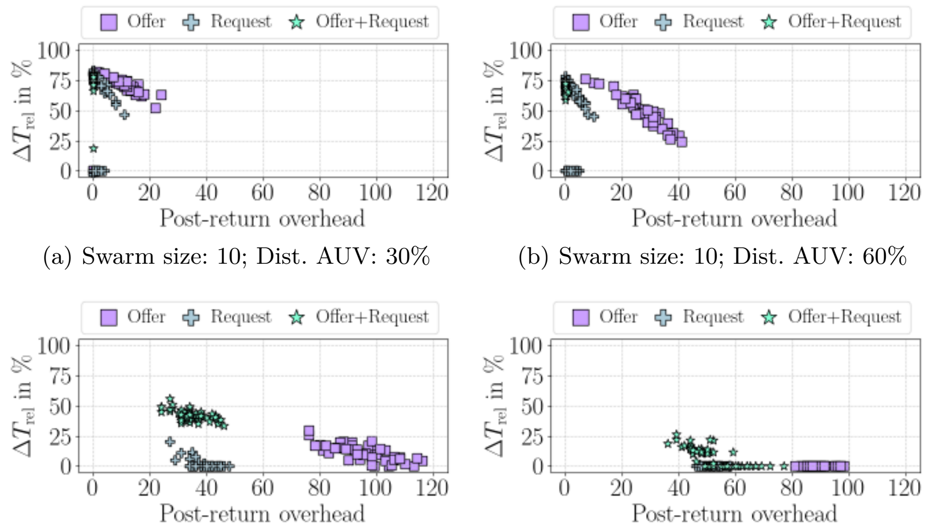

(a) Swarm size: 10; Dist. AUV: 30%

(b) Swarm size: 10; Dist. AUV: 60%

(c) Swarm size: 30; Dist. AUV: 30%

(d) Swarm size: 30; Dist. AUV: 50%

Fig. 5. Correlation between the relative improvement in mission duration and the additional number of waypoints (including the base) across task redistribution strategies.

To evaluate coordination efficiency, we analyze the relative improvement in mission duration among different strategies. We use the metric

$$\Delta T_{\mathrm{rel}} = \frac{T_{\max} - T_i}{T_{\max}} \cdot 100\% \tag{10}$$

where $T_{\max}$ is the non-cooperative baseline mission duration measured when agents operate independently without any communication or coordination. This baseline represents the worst-case scenario, i.e., the longest mission time. Improvements in coordination are reflected by higher values of ΔT_{rel}. In our setup, agents that finish their assigned waypoints return to the base and enter standby mode. This approach limits unnecessary travel and provides a common meeting point for task redistribution. Idle agents remain available to quickly take over leftover tasks, which is especially beneficial for agents at the network periphery or those temporarily disconnected. However, this redeployment can increase total travel distance and mission time.

Figure 5 shows mission duration versus the number of post-return overhead – additional points caused by repeated returns to base. The plots indicate that the `Offer` strategy causes the most post-return overhead because help is only offered after an AUV completes its tasks, resulting in frequent returns to base for reassignment. The `Request` and combined `Offer+Request` strategies generate similar overhead in terms of additional waypoints visited. Nevertheless,

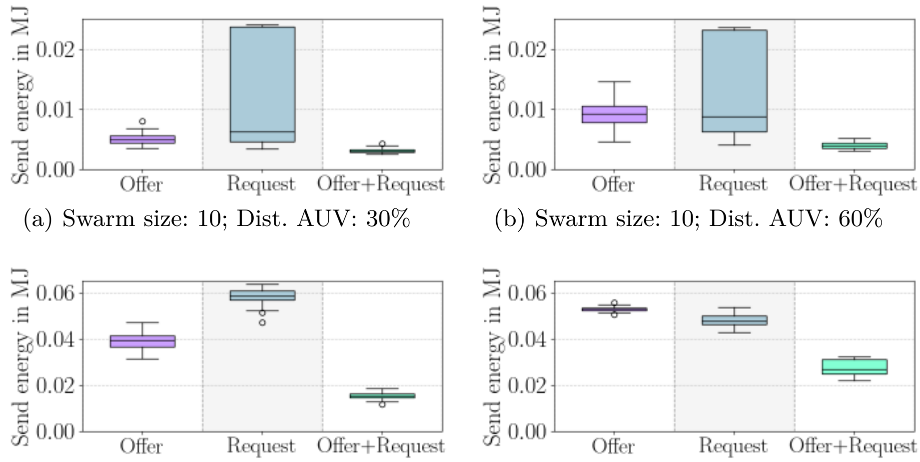

(a) Swarm size: 10; Dist. AUV: 30%

(b) Swarm size: 10; Dist. AUV: 60%

(c) Swarm size: 30; Dist. AUV: 30%

(d) Swarm size: 30; Dist. AUV: 50%

Fig. 6. Energy consumption for sending messages across the swarm.

the combined strategy significantly improves overall mission duration, reducing unnecessary redeployments and thus increasing coordination efficiency.

5.2 Communication Effort

AUVs typically use separate power supplies for propulsion and onboard electronics, including communication systems [11,13,2], as propulsion requires high power while communication demands low-noise, low-voltage sources. We therefore assess communication energy independently from propulsion.

Effective communication is essential for coordination, but it can also be energy-intensive and may hinder mission performance if not managed properly. Utilizing data from the *ahoi* acoustic modem [19,20] (which consumes 2.1 W for transmission and 0.3 W for reception, and sends 64-byte messages at a rate of 32.5 Bytes/s), we calculated the energy costs associated with each message. By combining this information with the transmission counts shown in fig. 6 and the reception logs in fig. 7, we derived total estimates for communication energy consumption.

The plots show that for all tested scenarios, the combined `Offer+Request` strategy achieves the lowest energy consumption for both sending and receiving messages. From an energetic perspective, it is therefore the most efficient strategy for communication. When comparing the individual strategies `Offer` and `Request`, distinct patterns emerge. In small swarms, the median energy consumption for communication is similar for both strategies; however, the standard deviation for `Request` is significantly higher, indicating less predictability in energy demand. The generally greater energy usage associated with `Request`

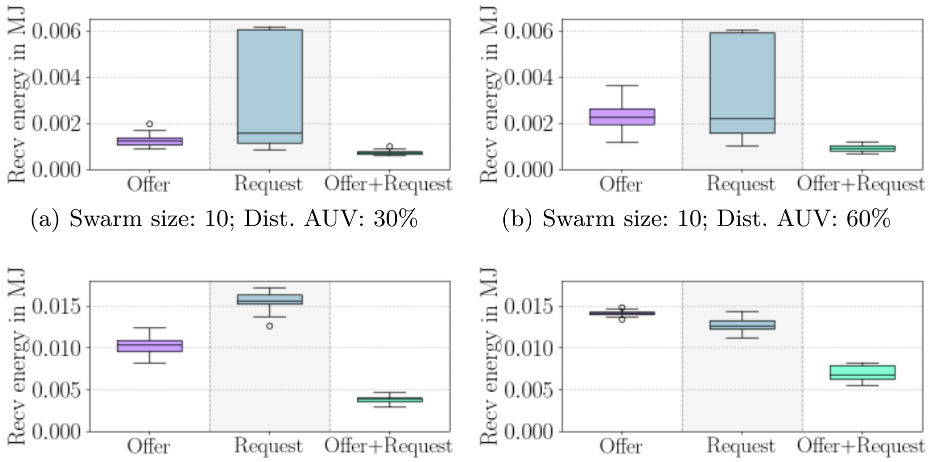

(a) Swarm size: 10; Dist. AUV: 30%

(b) Swarm size: 10; Dist. AUV: 60%

(c) Swarm size: 30; Dist. AUV: 30%

(d) Swarm size: 30; Dist. AUV: 50%

Fig. 7. Energy consumption for receiving messages across the swarm.

suggests that it involves more frequent responses, which implies increased connectivity between agents. This effect becomes more pronounced in large swarms, where the likelihood of isolated AUVs increases. Isolated agents send fewer messages because they receive fewer requests. In contrast, the `Offer` strategy is only triggered once an agent has completed all its tasks. While this reduces communication overhead when missions are progressing smoothly, the fixed TDMA slots can introduce unnecessary delays in information exchange. In some cases, this can lead to slower responses from other agents and hinder timely task redistribution.

5.3 Quantifying Swarm Opinion Formation

We evaluate collective decision-making using two key measures. The consensus score

$$C = \frac{1}{N} \sum_{i=0}^{N-1} \mathbb{I}\left(|x_i(k) - \bar{x}(k)| < \varepsilon\right), \tag{11}$$

quantifies how closely agents align their internal assessments, with $\bar{x}(k)$ as the swarm-wide average opinion and ε a small tolerance margin. The indicator function $\mathbb{I}(\cdot)$ returns 1 if agent i agrees with the majority within this margin. In contrast, Shannon entropy

$$H = -\sum_{\chi=1}^{\mathrm{X}} p_i \log_2 p_i. \tag{12}$$

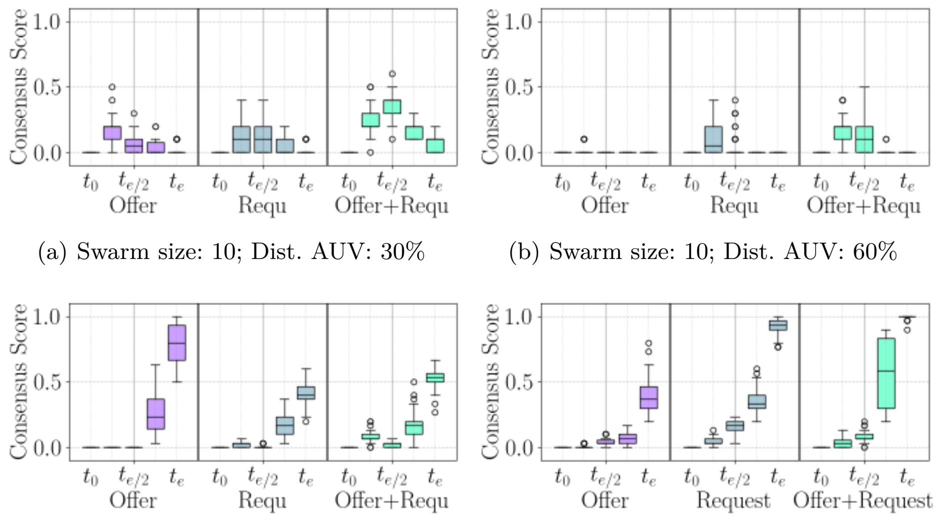

(a) Swarm size: 10; Dist. AUV: 30%

(b) Swarm size: 10; Dist. AUV: 60%

(c) Swarm size: 30; Dist. AUV: 30%

(d) Swarm size: 30; Dist. AUV: 50%

Fig. 8. Consensus score distribution at different mission phases. Higher values indicate stronger agreement within the swarm.

measures opinion diversity by binning normalized opinions $x_i(k) \in [0, 1]$ into χ intervals with corresponding frequencies p_i. Higher entropy indicates greater uncertainty or persistent disagreement. Together, these measures help differentiate genuine agreement from misleading uniformity and show whether the diversity arises from resilience or fragmentation. A high consensus score with low entropy suggests strong alignment among participants, while a low consensus score with high entropy indicates ongoing disagreement or indecision.

All strategies start with near-zero consensus in fig. 8. In small swarms with 30 % disturbed AUVs, `Offer` jumps quickly early and then plateaus, and `Request` remains unchanged mainly, `Offer+Request` shows a steady mid-mission rise in consensus before a slight decline. With 60% disturbed AUVs, `Offer` stays near zero, `Request` produces only isolated mid-mission peaks, and `Offer+Request` exhibits the same qualitative pattern as with 30% but with weaker alignment.

Entropy patterns mirror these trends: `Offer` shows a sharp early increase in entropy that returns to baseline, indicating quick but short-lived disagreement followed by fast convergence. `Request` reaches its peak mid-mission with high variability, reflecting scattered and less synchronized behavior. `Offer+Request` maintains moderate entropy on a plateau before decreasing, suggesting a balanced relationship between diversity and agreement that supports adaptive coordination.

For small swarms, the combined `Offer+Request` strategy performs best: it achieves meaningful mid-mission consensus while preserving moderate en-

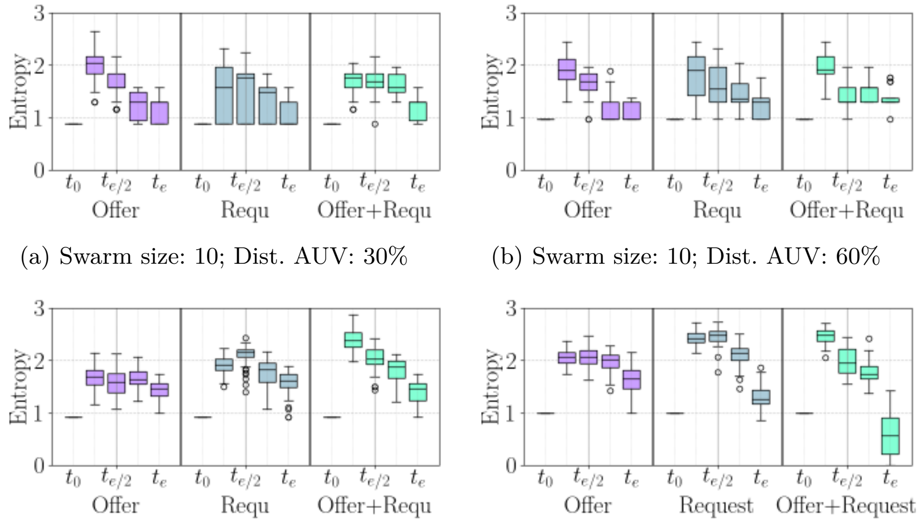

(a) Swarm size: 10; Dist. AUV: 30%

(b) Swarm size: 10; Dist. AUV: 60%

(c) Swarm size: 30; Dist. AUV: 30%

(d) Swarm size: 30; Dist. AUV: 50%

Fig. 9. Entropy of the opinion distribution at different mission phases. Higher values reflect greater disagreement and uncertainty among AUVs.

tropy, i.e., a helpful balance between agreement and diversity that supports adaptive, robust coordination. `Offer` yields fast initial alignment but sacrifices adaptability (resulting in a brittle consensus), and `Request` produces unpredictable, high-variance behavior with limited overall convergence. Thus, if the goal is reliable and adaptable task redistribution in well-connected small swarms, `Offer+Request` is the recommended strategy.

In large swarms, all strategies show a similar pattern: consensus rises only slightly at the start but increases sharply toward the end of the mission, resembling an exponential curve. The entropy trends confirm this behavior: it jumps early in the mission and then gradually decreases toward the end. This indicates that large swarms tend to form a collective assessment only late in the mission, likely when mission outcomes or resource shortages become apparent. In terms of mission success, this delayed consensus can have mixed effects. While it allows the swarm to remain flexible early on, it may also hinder timely task redistribution under high disturbance. Consequently, strategies that combine `Offer+Request` achieve higher completion rates in disturbed conditions, as they allow for both proactive offers and reactive requests, mitigating the risks of late consensus. The `Offer` strategy alone performs best when disturbances are low, but under high disturbance, its late consensus formation correlates with mission failure for many waypoints.

6 Conclusion and Outlook

This work investigates how leveraging agent stubbornness enhances cooperative task redistribution in communication-constrained AUV swarms, focusing on the decision-making process when unexpected events occur. Stubbornness acts as a tunable parameter balancing individual task commitment and collective flexibility. By integrating stubbornness into distributed opinion dynamics, AUVs can make autonomous decisions about when to cooperate, when to resist change, and how to interpret swarm dynamics, even under low-bandwidth, intermittent communication.

We evaluated three cooperative strategies. In the `Offer` strategy, agents proactively propose tasks once their own are complete; in `Request`, agents seek assistance for unfinished tasks; the combined `Offer+Request` strategy integrates both approaches. Across all scenarios, `Offer+Request` consistently achieved the most robust performance by balancing diversity and consensus, avoiding premature lock-in, and enabling effective task redistribution even in disturbed conditions. Small swarms coordinated effectively despite a substantial fraction of agents being less stubborn or disturbed, while large swarms tended to form consensus only late in the mission. This delayed consensus allows early flexibility but can hinder timely redistribution under high disturbance, again highlighting the value of moderate, context-sensitive stubbornness.

A key insight is that mission success correlates with well-timed, moderate consensus growth rather than maximal early agreement. Stubbornness mediates this by allowing agents to resist premature alignment while remaining responsive to critical task demands, effectively linking individual behavior to emergent swarm performance. Strategies that preserved diversity early and converged toward the end (such as `Offer+Request`) demonstrated the best combination of reliability, efficiency, and robustness. These findings reinforce prior insights from swarm intelligence and opinion dynamics [22,4,17], emphasizing that a balance between agreement and heterogeneity is vital for resilient group behavior in dynamic environments.

Future work should explore adaptive communication protocols that dynamically allocate bandwidth based on agent urgency or network topology, addressing the limitations of fixed TDMA slots and supporting timely task redistribution, in line with the principles of the FJ model [23]. Further investigation could also examine context-sensitive adjustment of stubbornness in real time, allowing agents to adapt their persistence to changing environmental conditions, mission priorities, or energy constraints. Extending the framework to heterogeneous swarms, mixed autonomy levels, or multi-objective missions may provide additional insights into scalable, resilient coordination strategies for real-world underwater operations.

References

1. Alfred-Wegener-Institut, Helmholtz-Zentrum für Polar- und Meeresforschung: Tiefsee-Ökologie und -Technologie (2024), https://www.awi.de/forschung/

biowissenschaften/tiefsee-oekologie-und-technologie/technologie.html, accessed: 27 Jan 2025

2. von Benzon, M., Sørensen, F.F., Uth, E., Jouffroy, J., Liniger, J., Pedersen, S.: An Open-Source Benchmark Simulator: Control of a Bluerov2 Underwater Robot. Journal of Marine Science and Engineering **10**(12), 1898 (2022)
3. Connor, J., Champion, B., Joordens, M.A.: Current Algorithms, Communication Methods and Designs for Underwater Swarm Robotics: A Review. IEEE Sensors Journal **21**(1), 153–169 (2020)
4. Couzin, I.D., Krause, J., Franks, N.R., Levin, S.A.: Effective Leadership and Decision-Making in Animal Groups on the Move. Nature **433**(7025), 513–516 (2005)
5. DeGroot, M.H.: Reaching a Consensus. Journal of the American Statistical Association **69**(345), 118–121 (1974)
6. Fang, Y., Huang, Z., Pu, J., Zhang, J.: AUV Position Tracking and Trajectory Control Based on Fast-Deployed Deep Reinforcement Learning Method. Ocean Engineering **245**, 110452 (2022)
7. Faruque, S.: Time Division Multiple Access (TDMA). Radio Frequency Multiple Access Techniques Made Easy pp. 35–43 (2019)
8. Frenkel, W., Renner, B.C.: Towards Energy-Aware Path Planning for AUV Swarms. In: Proceedings of the 18th International Conference on Control, Automation, Robotics and Vision (ICARCV 2024). pp. 1106–1111. IEEE, Dubai, United Arab Emirates (2024)
9. Frenkel, W., Renner, B.C.: Hyperparameter Optimization for PSO-Based Energy-Aware Path Planning for AUV Swarms. In: 2025 38th International Conference on Architecture of Computing Systems (ARCS). p. in print. Springer (2025)
10. Hadi, B., Khosravi, A., Sarhadi, P.: Deep Reinforcement Learning for Adaptive Path Planning and Control of an Autonomous Underwater Vehicle. Applied Ocean Research **129**, 103326 (2022)
11. Hasabnis, P., Pawar, P., Kulkarni, R., Sonawane, D., Mehta, C., Parseja, A., Pardeshi, A., Patil, J., Dandapur, A., Dandekar, H., et al.: Ekagra, an Autonomous Underwater Vehicle: Cost-Effective System Design and Implementation. In: OCEANS 2022-Chennai. pp. 1–7. IEEE (2022)
12. Heidemann, J., Stojanovic, M., Zorzi, M.: Underwater Sensor Networks: Applications, Advances and Challenges. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **370**(1958), 158–175 (2012)
13. Herrin, B., Close, V., Berner, N., Hebert, J., Reussow, E., James, R., Woodward, C., Mindlin, J., Paez, S., Bretas, N., et al.: Modularis: Modular Underwater Robot for Rapid Development and Validation of Autonomous Systems. In: OCEANS 2023-MTS/IEEE US Gulf Coast. pp. 1–7. IEEE (2023)
14. Luvisutto, A., Celani, A., Renda, F., Stefanini, C., De Masi, G.: Enhancing Collaboration in Uncertain Environment: Multi-Agent Reinforcement Learning for Underwater Monitoring. Expert Systems with Applications **277**, 127256 (2025)
15. Milot, A., Chauveau, E., Lacroix, S., Lesire, C.: Market-Based Multi-Robot Coordination with HTN Planning. In: 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 2606–2612. IEEE (2021)
16. Morozs, N., Mitchell, P., Zakharov, Y.V.: TDA-MAC: TDMA without Clock Synchronization in Underwater Acoustic Networks. IEEE Access **6**, 1091–1108 (2017)
17. Mäs, M., Flache, A., Helbing, D.: Individualization as Driving Force of Clustering Phenomena in Humans. PLoS computational biology **6**(10), e1000959 (2010)
18. Proskurnikov, A.V., Tempo, R.: A Tutorial on Modeling and Analysis of Dynamic Social Networks. Part I. Annual Reviews in Control **43**, 65–79 (2017)

19. Renner, B.C.: ahoi: Open-Source, Inexpensive, Low-Power Underwater Acoustic Modem (2023), `https://www.ahoi-modem.de`, accessed: 19. Dec 2024
20. Renner, B.C., Heitmann, J., Steinmetz, F.: AHOI: Inexpensive, Low-Power Communication and Localization for Underwater Sensor Networks and μAUVs. ACM Transactions on Sensor Networks (TOSN) **16**(2), 1–46 (2020)
21. Reuter, J., Mank, E., Aschemann, H., Rauh, A.: Battery State Observation and Condition Monitoring Using Online Minimization. In: 2016 21st International Conference on Methods and Models in Automation and Robotics (MMAR). pp. 1223–1228. IEEE (2016)
22. Surowiecki, J.: The Wisdom of Crowds. NY.: Anchor (2005)
23. Wang, L., Xing, Y., Johansson, K.H.: On Final Opinions of the Friedkin-Johnsen Model Over Random Graphs with Partially Stubborn Community. In: 2024 IEEE 63rd Conference on Decision and Control (CDC). pp. 4562–4567. IEEE (2024)
24. Wang, Z., Li, J., Li, J., Liu, C.: A Decentralized Decision-Making Algorithm of UAV Swarm with Information Fusion Strategy. Expert Systems with Applications **237**, 121444 (2024)
25. Wasserman, S.: Social Network Analysis: Methods and Applications. The Press Syndicate of the University of Cambridge (1994)
26. Xing, Y., Johansson, K.H.: A Concentration Phenomenon in a Gossip Interaction Model with Two Communities. In: 2022 European Control Conference (ECC). pp. 1126–1131. IEEE (2022)
27. Xing, Y., Johansson, K.H.: Concentration in Gossip Opinion Dynamics Over Random Graphs. SIAM Journal on Control and Optimization **62**(3), 1521–1545 (2024)
28. Zhang, C., Cheng, P., Du, B., Dong, B., Zhang, W.: AUV Path Tracking with Real-Time Obstacle Avoidance via Reinforcement Learning Under Adaptive Constraints. Ocean Engineering **256**, 111453 (2022)