Multimodal-Aware Meta-Adaptive Cross-Layer MAC for Fixed-Mobile Underwater Acoustic Networks

Huajie Zhang
School of Computer Science and Technology, Hainan University
Innovation Platform for Academicians of Hainan Province
Haikou, China
Email: zhhj@hainanu.edu.cn

Rongxin Zhu
School of Computer Science and Technology, Hainan University
Innovation Platform for Academicians of Hainan Province
Haikou, China
Email: zhurongxin@hainanu.edu.cn

Yupeng Ma
School of Computer Science and Technology, Hainan University
Innovation Platform for Academicians of Hainan Province
Haikou, China
Email: phpl@hainanu.edu.cn

Pengcheng Li
School of Computer Science and Technology, Hainan University
Innovation Platform for Academicians of Hainan Province
Haikou, China
Email: lpc1997@hainanu.edu.cn

Azzedine Boukerche
EECS
University of Ottawa
Ottawa, Canada
Email: boukerch@site.uottawa.ca

Qiuling Yang
School of Computer Science and Technology, Hainan University
Innovation Platform for Academicians of Hainan Province
Haikou, China
Email: qlyang@hainanu.edu.cn

***Abstract*—Fixed-mobile underwater acoustic networks must coordinate low-power monitoring nodes with autonomous underwater vehicles (AUVs) and gliders over links that have narrow bandwidth, long propagation delay, intermittent contact, and non-stationary sea-state dynamics. Fixed access rules are therefore inefficient: a slot that appears idle may expire before feedback returns, and parameters selected under one sea state may become unsuitable after a shift. This paper presents a multimodal-aware and meta-adaptive cross-layer medium access control (MAC) protocol for fixed-mobile underwater acoustic networks. The protocol fuses historical channel state information, environmental observations, and MAC telemetry through a lightweight temporal convolutional network (TCN) to produce slot-availability priors. Mobile nodes then construct short-validity Medium Access Plan (MAP) schedules under contact-window, fairness, retransmission, and expiration constraints. Feedback-triggered lightweight adjustment refines the transmission-power and contention-window hints within the MAP boundary. Simulations show that the predictor achieves accurate received signal strength indicator (RSSI) estimation ($R^2 = 0.8253$), the dynamic MAP scheduler improves relative delivery robustness, fairness, and energy efficiency under fixed-mobile-resource overload settings, and the supporting adjustment improves post-shift access behavior in the tested scenarios.**

***Index Terms*—Underwater acoustic networks, fixed-mobile collaboration, cross-layer MAC, multimodal prediction, dynamic MAP scheduling, meta-learning.**

This work was partially supported by the National Natural Science Foundation of China under Grant 62662027 and 62362026; in part by the Key Project of Hainan Province under Grant ZDYF2023GXJS158; in part by Hainan Provincial Natural Science Foundation of China under grant 626MS0092 and 626QN0556. (Corresponding author: Rongxin Zhu)

I. Introduction

Underwater acoustic networks enable ocean observation, environmental monitoring, seabed exploration, and cooperative underwater operations. Their medium access control (MAC) design is fundamentally different from terrestrial radio networks because acoustic links have narrow bandwidth, low propagation speed, severe multipath, large delay spread, Doppler variation, and strong environmental noise. These physical characteristics make carrier sensing unreliable, enlarge guard-time requirements, and reduce the effectiveness of fixed-slot or centrally coordinated access unless the MAC layer is redesigned for acoustic propagation [1], [2].

Fixed-mobile underwater networks add another layer of difficulty. Fixed monitoring nodes are suitable for long-term sensing but have limited energy and local visibility. Autonomous underwater vehicles (AUVs) and gliders can collect data, estimate local channel conditions, and coordinate nearby fixed nodes, but their contact windows are short and topology-dependent. Recent underwater MAC protocols improve reliability through dynamic reservation, hybrid retransmission, non-orthogonal multiple access (NOMA)-assisted access, georouting-aware control, AUV-assisted scheduling, cluster-based spatial–temporal scheduling, and mobility-aware conflict-graph scheduling [3]–[9]. AUV-assisted path repair has also been studied to restore disrupted underwater routing paths [10]. Topology-aware node scheduling and opportunistic forwarding have also been investigated for underwater networks [11], [12]. These studies improve channel utilization,

979-8-3195-0167-7/26/$31.00 ©2026 IEEE

but most of them still react to current requests or current channel estimates. A practical fixed-mobile MAC should also answer whether a future contact interval will remain usable after propagation delay, data transmission, guard time, and feedback are considered.

Learning-based underwater communication methods provide useful tools for channel prediction and adaptive control [13]–[16]. However, prediction, scheduling, and access-parameter adaptation are often optimized as separate modules. This separation is fragile under non-stationary sea states: after a sudden increase in noise or delay spread, the same power and contention-window settings may produce acknowledgment (ACK) timeouts, power-insufficient transmissions, and unnecessary backoff [17]. Meta-learning is attractive in this setting because it can learn an initialization that adapts with only a small support set [18], while drift detection can decide when such online updates are necessary [19].

This paper proposes MAMA-MAC, a compact cross-layer protocol that connects multimodal prediction, mobile-node-driven Medium Access Plan (MAP) scheduling, and feedback-triggered access adjustment. Here, *multimodal* refers to heterogeneous data modalities—channel observations, environmental observations, and MAC telemetry—rather than multiple physical-layer communication technologies. The central design principle is to use learning as a soft prior and preserve acoustic feasibility as a hard MAC-layer constraint. In particular, predicted slot availability increases the scheduling score, but residual contact time, MAP expiration, retransmission constraints, and fairness penalties decide whether a transmission is actually authorized. The two contributions are summarized as follows:

- We design a closed-loop cross-layer MAC framework that couples a TCN-based slot-availability prior with mobile-node-driven dynamic MAP scheduling. The scheduler retains contact-window, expiration, retransmission, and fairness rules as hard constraints.
- We incorporate a drift-triggered, lightweight access-adjustment path that uses feedback to refine power and contention-window hints in later MAP cycles. This supporting mechanism is explicitly bounded by the MAP schedule and does not authorize transmissions by itself.

The remainder of this paper is organized as follows. Section II defines the fixed-mobile network model and problem statement. Section III presents MAMA-MAC. Section IV reports simulation results, and Section V concludes the paper.

II. System Model and Problem Statement

Consider a fixed-mobile underwater acoustic network with fixed monitoring nodes $\mathcal{S} = \{s_1, \ldots, s_N\}$ and mobile nodes $\mathcal{A} = \{a_1, \ldots, a_M\}$. Fixed nodes generate sensing packets and send compact requests (REQ) messages when they have pending traffic. Mobile nodes patrol the area, collect requests from reachable fixed nodes, broadcast dynamic MAP schedules, receive data, and return ACK-type feedback. One scheduling cycle therefore contains request collection, multimodal prediction, MAP construction, scheduled data transmission, feedback collection, and optional adaptation.

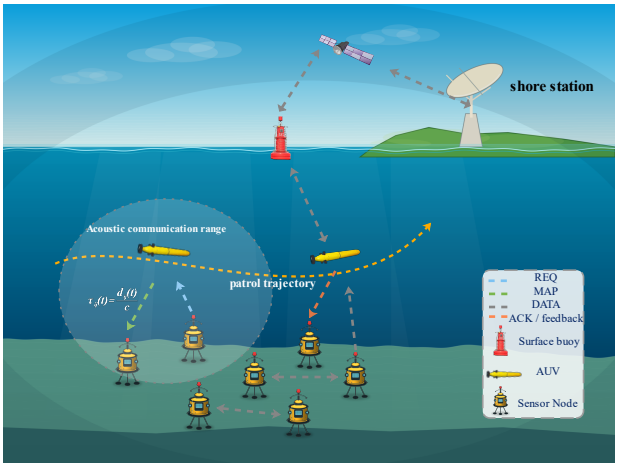


Fig. 1. Fixed-mobile underwater acoustic network with REQ, MAP, DATA, ACK, synchronization, and environmental-observation flows.

Fig. 1 shows the considered fixed-mobile model. Fixed nodes send requests and data to nearby mobile collectors, while mobile nodes return MAP and ACK feedback, synchronize regional schedules, and use environmental observations as cross-layer context. The figure emphasizes that connectivity is opportunistic: a fixed node is schedulable only while a mobile node remains within a usable acoustic contact window.

In this model, fixed nodes are energy-constrained and remain mostly silent unless sensing data are pending. Mobile nodes have stronger sensing, storage, and scheduling capabilities, and therefore act as local coordinators during their contact periods. The proposed protocol does not require a permanently connected sink or a global time-slot table. Instead, each mobile node constructs short-validity MAPs from the requests and channel observations available in its local service region. This design is suitable for sparse underwater deployments where long propagation delay and intermittent contact make frequent centralized coordination costly.

For link (s_i, a_j) at time t, the distance is $d_{ij}(t)$ and the sound speed is c. The one-way propagation delay is

$$\tau_{ij}(t) = \frac{d_{ij}(t)}{c}. \tag{1}$$

In the simulations, c is fixed to the nominal value $1500\,\mathrm{m/s}$. In practical deployments, sound speed varies with environmental conditions, so an estimation error in c introduces a corresponding error in $\tau_{ij}(t)$. For a small perturbation Δc, the relative delay error is approximately $\Delta\tau/\tau \approx -\Delta c/c$. The present scheduler therefore retains a guard-time margin and feedback-based delay telemetry, while explicit online sound-speed estimation is outside the scope of this evaluation.

A flow is schedulable only when the link margin is sufficient and the residual contact window $T^c_{ij}(t)$ can cover data transmission, guard time, and feedback. The acoustic attenuation, environmental noise, sea state, and relative motion

are abstracted into features such as link gain, signal-to-noise ratio (SNR) margin, propagation delay, and delay spread. Although delay spread is fundamentally a physical-layer channel characteristic, it is used here only as a cross-layer indicator of multipath severity and potential delivery uncertainty, rather than as a physical-layer control variable. These features are observable by mobile nodes or estimated from recent feedback. This feedback-based model follows recent underwater link-adaptation studies [20].

The cross-layer observation at time t is

$$\mathbf{x}_t = [\mathbf{c}_t, \mathbf{e}_t, \mathbf{m}_t], \tag{2}$$

where $\mathbf{c}_t$ contains the received signal strength indicator (RSSI), SNR, and root-mean-square (RMS) delay spread; $\mathbf{e}_t$ contains sea state, environmental noise, wind speed, water temperature, salinity, depth, propagation delay, link gain, and relative velocity; and $\mathbf{m}_t$ contains queue length, retransmission count, recent success rate, collision rate, timeout rate, ACK round-trip-time (RTT) exponentially weighted moving average (EWMA), and queue-delay EWMA.

For a mobile node a_j in scheduling round r, let $\mathcal{F}_{j,r}$ be the candidate flow set. Each flow includes priority, data length, queue length, retransmission history, predicted availability, and contact-window information. The scheduling objective is to assign feasible non-overlapping intervals and access hints that increase expected delivery and service fairness while bounding delay, energy cost, repeated scheduling, and control overhead.

The design follows three constraints. First, the acoustic control channel is expensive; therefore, one MAP should authorize several transmissions whenever possible instead of requiring per-packet handshakes. Second, prediction uncertainty must not directly create unsafe transmissions. A high predicted access probability only raises the scheduling score, while contact-window feasibility and MAP expiration remain hard constraints. Third, adaptation must be event-triggered. Continuous online learning would consume samples, energy, and time under stable sea states; thus, few-shot updates are invoked only when the feedback stream indicates a distribution shift.

A complete protocol cycle proceeds as follows. First, fixed nodes with pending packets transmit compact REQ messages containing flow identity, queue status, data length, and retry information. Second, the mobile node aggregates recent channel state information (CSI), environmental observations, and MAC telemetry into a temporal feature window and invokes the TCN predictor to obtain access priors for upcoming slots. Third, infeasible flows are removed according to contact-window and MAP-expiration constraints, while the remaining flows are ranked and inserted into the MAP. Fourth, fixed nodes transmit data only within their assigned MAP intervals and wait for ACK-type feedback. Finally, collected feedback updates MAC telemetry; when a persistent change is detected, it can refine the parameter hints used in a later MAP cycle.

The modules exchange compact information but do not override each other. The TCN output is used only as a soft scheduling prior. MAP timing, contact-window feasibility, expiration, and non-overlap constraints remain hard MAC-layer rules. The feedback-triggered adjustment only refines the power ratio and contention window used inside an assigned MAP interval; it cannot extend the scheduled duration, bypass expiration, or authorize an unscheduled transmission.

TABLE I
MAIN MESSAGE FIELDS IN ONE MAP CYCLE.

Message	Main fields	Role
REQ	Flow ID, queue length, priority, data bits, retry count, node position	Expose traffic demand and local state.
MAP	txStart, txDuration, expireTime, powerHint, cwHint, burst ID	Assign a short-validity access contract.
Feedback	ACK, timeout, collision, power fail, RTT, queue delay, margins	Update telemetry and bounded adjustment samples.

III. MAMA-MAC PROTOCOL

A. Multimodal Slot Prediction

The predictor takes an observation window $\mathbf{X}_{t-L+1:t}$ and estimates future channel states and slot availability. A causal, dilated TCN enlarges the temporal receptive field without recurrent state, allowing it to capture longer-term channel variation with limited online computation. Before inference, CSI, environmental observations, and MAC telemetry are aligned in the same temporal window. The auxiliary channel-state head predicts RSSI, SNR, and RMS delay spread, while the slot head produces the prior $\hat{p}_{t+k}$ for a future access opportunity.

For training, a sample is labeled available only when its SNR margin, contact time, and collision-risk proxy satisfy the preset MAC-feasibility thresholds. The loss combines slot-classification and channel-regression terms. The resulting availability prior is used to rank otherwise feasible flows; it never directly authorizes a transmission. The predictor therefore produces a MAC prior rather than a final decision. Hard constraints remain with the scheduler.

The TCN is trained offline and kept fixed during each online run. It provides a MAC prior, while hard constraints are enforced by the scheduler and feedback-triggered adjustment is handled by the lightweight access controller.

B. Mobile-Node-Driven MAP Scheduling

In this paper, MAP denotes Medium Access Plan, i.e., a MAC-layer access schedule broadcast by the mobile node, rather than maximum a posteriori estimation.

During the request collection window, each candidate flow is represented by its source node, selected mobile receiver, data length, queue state, priority, creation time, retry count, and recent scheduling status. After scoring and feasibility filtering, the mobile node writes the selected entries into the MAP. A fixed node transmits only when its flow identifier matches a valid MAP entry and the current time is within the assigned interval before expiration. The fields `txStart` and `txDuration` define the access window, while `powerHint`

and `cwHint` provide preferred access parameters for the scheduled transmission.

After collecting requests, a mobile node scores each candidate flow f by combining traffic demand, prediction confidence, urgency, fairness, and access cost:

$$\begin{aligned} S(f) = {} & w_p P(f) + w_{\hat{p}} \hat{p}(f) + w_q Q(f) + w_d U_d(f) \\ & + w_w U_w(f) + w_c U_c(f) + w_r R(f) + w_f F(f) \\ & - w_l L(f) - w_{rp} \Pi_f(f) - w_{sp} \Pi_s(f). \end{aligned} \tag{3}$$

Here $P(f)$ is traffic priority, $\hat{p}(f)$ is the predicted slot-availability prior, $Q(f)$ is queue length, $U_d(f)$, $U_w(f)$, and $U_c(f)$ denote delay, waiting-time, and contact-window urgency, $R(f)$ is retransmission urgency, $F(f)$ is fairness compensation, and $\Pi_f(f)$ and $\Pi_s(f)$ are flow-level and node-level repeated-scheduling penalties. The fairness and cost terms follow recent underwater MAC studies that jointly consider reliability, delay, energy, and fairness [21], [22]. The normalized transmission cost is

$$L(f) = \frac{T_{\mathrm{data}}(f) + T_{\mathrm{guard}}}{T_{\mathrm{frame}}}. \tag{4}$$

A candidate is retained only if

$$T_c(f) \geq T_{\mathrm{data}}(f) + T_{\mathrm{guard}}. \tag{5}$$

The final MAP entry must also leave enough time for data transmission, ACK feedback, two-way propagation, guard time, contact validity, and MAP expiration:

$$\begin{aligned} & t_s(f) + T_{\mathrm{data}}(f) + T_{\mathrm{ACK}} + 2\tau(f) + T_{\mathrm{guard}} \\ & \leq \min\{t_{\mathrm{contact}}(f), t_{\mathrm{exp}}(f)\}. \end{aligned} \tag{6}$$

This condition means that the data leg, ACK leg, propagation delay, guard time, contact deadline, and MAP expiration are jointly feasible. Assigned MAP entries are non-overlapping. This delay-aware constraint is consistent with recent learning-based underwater MAC designs [23]. With

$$t_e(f) = t_s(f) + T_{\mathrm{data}}(f) + T_{\mathrm{guard}}, \tag{7}$$

the scheduler enforces

$$[t_s(f_i), t_e(f_i)) \cap [t_s(f_j), t_e(f_j)) = \emptyset, \quad \forall f_i \neq f_j. \tag{8}$$

Burst scheduling is enabled only when the first flow from the same source has sufficient delivery confidence. If recent failures, low link confidence, or short residual contact time are observed, the burst length is limited to avoid wasting the contact window. The scheduler also applies repeated-scheduling penalties at both the flow level and the node level. The former suppresses repeated insertion of the same unfinished flow into consecutive MAPs, while the latter prevents one temporarily favorable node from occupying multiple scheduling cycles. High-priority flows may receive a cooldown exemption, but they are still constrained by contact-window, expiration, and non-overlap rules.

This priority rule lets waiting and retransmission urgency lift moderate-prior flows when the contact window is safe, while rejecting high-prior flows that violate MAP boundaries or fairness penalties.

C. Feedback-Triggered Lightweight Access Adjustment

When sea state changes, static access parameters may become invalid. The access action is

$$a = (\rho, CW), \tag{9}$$

where ρ is the transmission power ratio and CW is the contention window. The finite access-action space is

$$\mathcal{A}_p = \{\rho_1, \ldots, \rho_P\} \times \{CW_1, \ldots, CW_C\}. \tag{10}$$

The policy uses a compact cross-layer feature vector including link margin, delay spread, propagation delay, queue length, retry count, recent success rate, collision rate, timeout rate, and energy budget when available, and outputs action probabilities as

$$\pi_\theta(a \mid \mathbf{o}, \mathbf{m}) = \mathrm{softmax}(g_\theta(\mathbf{o}, \mathbf{m})). \tag{11}$$

The reward balances success, energy, delay, and failure:

$$R = \mathbb{I}_{\mathrm{success}} - \lambda_e E - \lambda_d D - \lambda_f \Phi. \tag{12}$$

Here, *meta-adaptive* refers to few-shot adjustment from an initialization learned across sea-state tasks during offline meta-training; it does not claim a new meta-learning algorithm. The Page–Hinkley detector only determines when adjustment is activated, whereas Reptile-lite provides the adaptation-oriented initialization. Offline Reptile-lite training follows a first-order initialization-based strategy [24]. Across multiple sea-state tasks, the method moves the shared initialization toward task-adjusted parameters:

$$\theta \leftarrow \theta + \beta(\theta_i' - \theta). \tag{13}$$

Page–Hinkley is a lightweight sequential change-detection method that accumulates persistent deviations of an observed statistic from its running reference level. Online, we use the Page–Hinkley statistic only as a gate for online adaptation by monitoring a score built from recent success rate, ACK RTT, queue delay, power margin, and SNR margin:

$$PH_t = \max(0, PH_{t-1} + z_t - \bar{z}_t - \delta). \tag{14}$$

The scalar score z_t is obtained by normalizing and linearly combining recent success, delay, power-margin, and SNR-margin indicators. When $PH_t > \eta$, the controller collects K support samples $\mathcal{D}_s = \{(\mathbf{o}_i, \mathbf{m}_i, a_i, r_i)\}_{i=1}^{K}$, updates a policy bias or lightweight head, and commits the adjusted ρ and CW to later MAP cycles. This lightweight adjustment is motivated by recent meta-learning-based underwater link adaptation [25]. When power-insufficient feedback dominates the support set, the controller increases the preference for higher-power actions. When collision feedback dominates, it increases the preference for larger contention windows. If high-power actions do not improve recent rewards, further power increase is capped to avoid unnecessary energy cost.

The detector follows four states: warm-up, monitor, cooldown, and stabilizing. Warm-up estimates the baseline, monitor accumulates deviation, cooldown avoids repeated triggers from the same disturbance, and stabilizing waits for the

TABLE II
COMPACT EVALUATION SETTINGS

Item	Setting
CSI dataset	Watermark benchmark, NCS1 channel [26]
CSI split/window	Trace-level 80/20 split, 50-step window
Simulator	OMNeT++
Area/sound speed	$3000\,\mathrm{m} \times 3000\,\mathrm{m}$, $1500\,\mathrm{m/s}$
Acoustic data rate	2000 bps
Mobile resources	5 AUVs, fixed sensors $N = 10$–40
Communication range	1200 m for AUVs, 800 m for sensors
Control packet size	REQ/ACK 128 bits; MAP 128 + 224 bits/entry
Experiment duration	MAC 1000 s; ablation 600 s
Repeated seeds	5 seeds for MAC and ablation runs
Baselines/variants	QueueContact, SlottedFAMA, StaticTDMA; Full variants



Fig. 2. Training and test loss curves of the TCN-based CSI predictor.

drift score and recent success rate to settle. Only confirmed triggers open the few-shot update path; the committed action preference is then applied in later MAP cycles.

ACK, timeout, collision, and power-fail outcomes update EWMA telemetry, including recent success, RTT, queue delay, and link margins. The updated telemetry is reused by the next prediction, MAP scheduling, and adaptation cycle, forming the closed cross-layer feedback loop.

The adjustment path is intentionally separated from MAP construction. The detector may request an update only after persistent feedback deviation, and the update changes the preference among a small, preconfigured set of power-ratio and contention-window actions. It cannot create a new entry, lengthen an assigned interval, or change an entry's expiration time. Therefore, a delayed or inaccurate drift trigger can affect the efficiency of later access attempts, but it cannot invalidate the contact-window and non-overlap guarantees enforced by the scheduler.

D. Complexity and Overhead

If a mobile node receives F candidate requests and can assign at most K_s slots, scoring is $O(F)$ and score sorting is $O(F \log F)$. MAP construction is therefore approximately $O(F \log F + K_s)$. TCN inference is bounded by the input length, channels, and convolutional layers, while online adaptation is event-triggered and uses a small support set. If each MAP entry requires B_e bits and the MAP header requires B_0 bits, the MAP overhead is

$$B_{\mathrm{MAP}} = B_0 + K_s B_e. \tag{15}$$

The total control overhead in one scheduling cycle can be summarized as

$$B_{\mathrm{ctrl}} = N_{\mathrm{REQ}} B_{\mathrm{REQ}} + B_{\mathrm{MAP}} + N_{\mathrm{ACK}} B_{\mathrm{ACK}} + B_{\mathrm{sync}}. \tag{16}$$

Because one MAP can authorize multiple DATA transmissions and short bursts, the amortized control cost per delivered packet decreases when more MAP entries are successfully used.

IV. EVALUATION

The evaluation focuses on three complementary aspects: CSI-based prediction and availability proxy, high-density MAC baseline comparison, and closed-loop ablation under sea-state shifts. The public CSI dataset validates channel-state prediction and a CSI-derived availability proxy, while the complete MAP scheduling constraints and feedback-triggered adjustment are evaluated in OMNeT++ simulations.

A. Experimental Settings

The CSI prediction experiment uses the public Watermark benchmark's NCS1 channel [26], which replays time-varying impulse responses measured on the Norwegian continental shelf. We derive RSSI, SNR, and RMS delay spread from the channel traces, use a trace-level 80/20 train/test split, and apply a five-layer TCN with a 50-step window to predict next-step channel indicators. Watermark does not provide synchronized environmental or MAC telemetry; these auxiliary features are supplied by the scenario and feedback models only in the OMNeT++ system experiments.

The MAC and feedback-adjustment experiments are implemented in OMNeT++. The high-density MAC comparison fixes the mobile resources to five AUVs and increases the number of fixed sensors from 10 to 40. The MAC baselines are `QueueContact`, `SlottedFAMA`, and `StaticTDMA`. The closed-loop ablation compares Full, NoPrediction, NoMeta, and NoDriftDetector under two sea-state shift scenarios. Unless otherwise stated, results for the MAC comparison and closed-loop ablation are averaged over five repeated seeds. The main metrics are success rate, throughput, delay, fairness, energy per delivered bit, post-shift success, and adaptation timing.

All MAC methods use the same deployment area, acoustic data rate, control-packet sizes, sensor density, and mobile-node budget listed in Table II. The comparison therefore isolates the access and scheduling decisions rather than increasing the resources available to the proposed method. The Watermark traces are used only for the CSI-prediction component and its availability proxy; the MAC results are generated by the trace-driven OMNeT++ scenarios, where contact windows, queued requests, collisions, and feedback can be controlled consistently across methods.

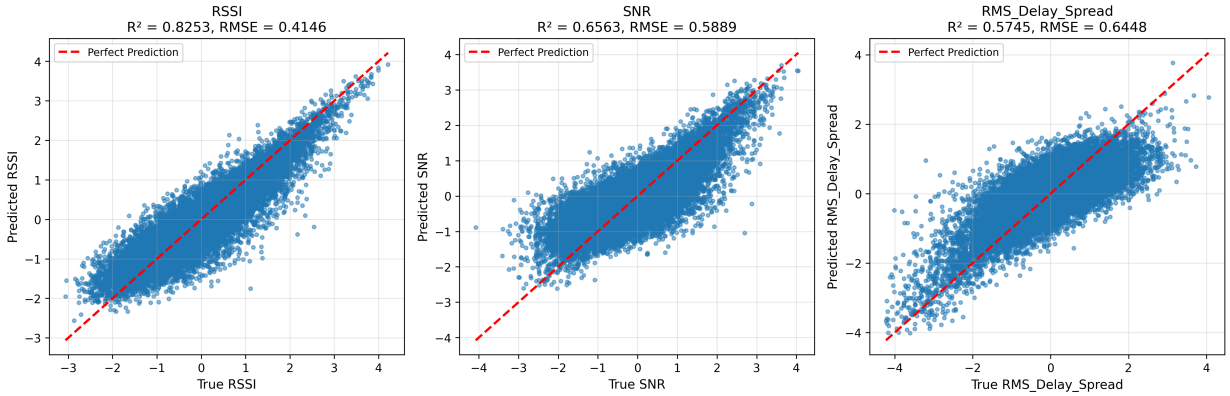
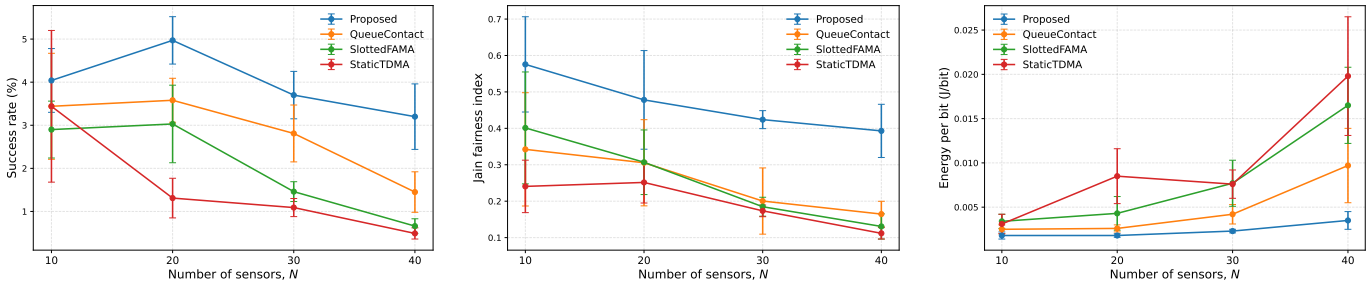


Fig. 3. Predicted and true values of RSSI, SNR, and RMS delay spread.



(a) Success rate (b) Fairness (c) Energy per bit

Fig. 4. High-density MAC baseline comparison with fixed mobile resources.

For the ablation, Full uses all three components: the TCN prior, the drift gate, and the bounded few-shot adjustment. No-Prediction removes the prior from the scheduling score while retaining feasibility checks; NoMeta disables the feedback-triggered parameter update; and NoDriftDetector retains the update path but removes the persistent-change gate. These variants share the same traffic generation, MAP validity rules, and shift traces, so their differences are attributable to the removed component rather than to a changed network configuration.

B. CSI Prediction and Availability Proxy

Fig. 2 shows stable learning behavior: the training loss decreases from about 0.63 to about 0.25, while the test loss decreases from about 0.57 to about 0.31 and flattens after roughly 40 epochs. Fig. 3 compares predicted and true CSI values. RSSI is the most learnable target, with $R^2 = 0.8253$ and RMSE $= 0.4146$. SNR and RMS delay spread obtain lower but useful accuracy, with $R^2 = 0.6563$ and $R^2 = 0.5745$, respectively.

The CSI-derived availability proxy evaluates whether the predictor provides a useful scheduling prior. A test sample is labeled available when the measured CSI satisfies SNR ≥ 4 dB, RSSI ≥ -2 dB, and RMSDelay ≤ 55. The resulting proxy classifier reaches F1 $= 0.7764$ and AUC $= 0.9031$ on 24,720 test sequences. This proxy does not claim to fully determine MAC slot availability. In the complete protocol, the OMNeT++ scheduler further combines the prediction prior with contact-window feasibility, collision-risk estimation, guard time, queue state, and feedback telemetry before admitting a flow into a MAP entry.

TABLE III
$N = 40$ OVERLOAD SUMMARY UNDER FIXED FIVE-AUV RESOURCES.

Method	Succ.	Fair.	E/bit (J/bit)	Sensor cov.
Proposed	0.0320	0.3930	0.0035	0.5450
QueueContact	0.0145	0.1646	0.0097	0.2700
SlottedFAMA	0.0066	0.1307	0.0165	0.1900
StaticTDMA	0.0049	0.1118	0.0198	0.1500

C. High-Density MAC Baseline Comparison

The high-density MAC experiment fixes the mobile resources to five AUVs and increases the fixed-sensor density from $N = 10$ to $N = 40$. This creates an overload stress setting in which more request flows compete for the same mobile service opportunities. Therefore, the results should be interpreted as relative robustness under overload rather than nominal deployment performance.

Fig. 4 compares the proposed scheduler with QueueContact, SlottedFAMA, and StaticTDMA. Under the fixed-5-AUV overload setting, all methods experience low absolute success rates as the sensor density increases, while the proposed scheduler maintains stronger relative robustness in success rate, fairness, and energy per delivered bit.

Fig. 4 shows the density trend, while Table III reports the exact $N = 40$ overload endpoint and adds sensor-level coverage. The endpoint confirms that MAP-based admission

TABLE IV
CLOSED-LOOP ABLATION SUMMARY UNDER TWO SHIFT SCENARIOS.

Scenario	Variant	SuccessRate	Throughput	Fairness	Energy/bit	Post-shift Success	Recovery Indicator	Adapt Commit
AbruptShift	Full	0.1012	268.80	0.5184	0.0031	0.0458	28.2	23
AbruptShift	NoPrediction	0.0935	249.67	0.5015	0.0034	0.0400	28.2	23
AbruptShift	NoMeta	0.0341	90.76	0.3624	0.0071	0.0000	40.0	N/A
AbruptShift	NoDriftDetector	0.0735	212.70	0.4786	0.0040	0.0325	15.0	31
LateHarsh	Full	0.1012	297.95	0.5070	0.0029	0.0211	25.6	34
LateHarsh	NoPrediction	0.1029	290.16	0.4946	0.0033	0.0244	25.2	34
LateHarsh	NoMeta	0.0606	179.25	0.4675	0.0040	0.0000	30.0	N/A
LateHarsh	NoDriftDetector	0.0924	270.29	0.5480	0.0030	0.0133	30.0	41

improves both packet-level delivery and node-level service coverage by filtering candidate flows before channel access.

At the $N = 40$ endpoint, the proposed method attains a success rate of 0.0320, compared with 0.0145 for the strongest listed baseline, `QueueContact`. It also raises fairness from 0.1646 to 0.3930, reduces energy per delivered bit from 0.0097 to 0.0035 J/bit, and increases sensor coverage from 0.2700 to 0.5450. These absolute success rates remain low because the setting deliberately overloads five mobile collectors; the useful comparison is that feasibility filtering and short-validity MAP admission preserve service for more sensors without increasing the available mobile resources.

D. Closed-Loop Ablation under Sea-State Shifts

To verify the full-system claim, we ablate prediction, meta adaptation, and drift detection on the same trace-driven simulator. The formal matrix contains four variants, two shift scenarios, and five repeated seeds, yielding 40 runs.

Table IV shows that the full closed loop provides the most stable trade-off among throughput, energy efficiency, post-shift behavior, and adaptation timing. Removing meta adaptation is the dominant failure mode: post-shift success becomes zero in both scenarios and the adaptation commit is unavailable. Removing drift detection preserves some traffic delivery before the controller reacts, but adaptation is delayed to sample 31 in AbruptShift and sample 41 in LateHarsh, compared with samples 23 and 34 for Full. Removing prediction weakens the scheduling prior and reduces throughput in AbruptShift. These results show that the gain is not produced by a single isolated module; CSI prediction, drift detection, and few-shot access adaptation need to be connected through feedback.

V. CONCLUSION

This paper presented MAMA-MAC, a multimodal-aware and meta-adaptive cross-layer medium access control (MAC) protocol for fixed-mobile underwater acoustic networks. The protocol fuses channel, environmental, and MAC telemetry features to predict access opportunities and lets mobile nodes generate dynamic Medium Access Plan (MAP) schedules under contact-window, expiration, retransmission, and fairness constraints. A feedback-triggered lightweight adjustment refines power and contention-window hints after persistent sea-state changes, while all transmission decisions remain bounded by the MAP schedule. The simulation results show that multimodal prediction provides useful scheduling priors, while dynamic MAP scheduling improves delivery performance, fairness, and energy efficiency under fixed-mobile-resource overload settings. The ablation results further show the value of the feedback path in the tested shift scenarios.

Extending the framework to mobile sensing nodes will require more frequent updates of contact windows, relative-motion estimates, and MAP-expiration information. Higher mobility may shorten MAP validity and increase scheduling overhead. Evaluating these effects, together with richer mobility traces, heavier traffic loads, real sea-trial channels, and stricter low-power deployment constraints, remains important future work.

REFERENCES

[1] R. Zhu, A. Boukerche, D. Li, and Q. Yang, "Delay-aware and reliable medium access control protocols for uwsns: Features, protocols, and classification," *Computer Networks*, vol. 252, p. 110631, 2024.

[2] A. S. Ghazy, N. Adam, I. C. Ijeh, O. Haddad, S. Karandikar, G. Kaddoum, C. Talhi, N. Iqbal, and H. S. Khallaf, "Underwater wireless communication systems and networks: A survey, future directions, and novel proposals," *IEEE Open Journal of the Communications Society*, 2026.

[3] H. Mei, H. Wang, X. Shen, Z. Jiang, W. Bai, C. Wang, and Q. Zhang, "An adaptive mac protocol for underwater acoustic sensor networks with dynamic-high load," *IEEE Sensors Journal*, vol. 24, no. 6, pp. 9059–9072, 2024.

[4] V.-P. Dao and T.-H. Tran, "Arq_hmac: An automatic repeat request strategy for hybrid mac protocol in underwater acoustic sensor networks," *IEEE Access*, vol. 13, pp. 159 196–159 206, 2025.

[5] X. Li, Z. Pu, W. Wang, C. Chi, and H. Huang, "Upnc-mac: A collision-aware and pd-noma-enabled mac protocol for distributed underwater acoustic networks," *IEEE Internet of Things Journal*, vol. 12, no. 23, pp. 51 125–51 141, 2025.

[6] B. Zheng, Z. Jiang, X. Shen, and W. Ning, "Gnc-mac: An efficient georouting-aware and cross-layer mac protocol based on noma for underwater acoustic networks," *IEEE Internet of Things Journal*, vol. 13, no. 10, pp. 21 424–21 439, 2026.

[7] J. Guo, S. Song, J. Liu, M. Pan, J.-H. Cui, and G. Han, "As-mac: An adaptive scheduling mac protocol for reducing the end-to-end delay in auv-assisted underwater acoustic networks," *IEEE Transactions on Mobile Computing*, vol. 24, no. 2, pp. 1197–1211, 2025.

[8] Q. Yang, Y. Chen, W. Dong, T. Li, R. Zhu, and X. Huang, "Cluster-based spatial–temporal mac scheduling protocol for underwater sensor networks," *IEEE Sensors Journal*, vol. 23, no. 15, pp. 17 690–17 702, 2023.

[9] R. Zhu, H. Zhang, P. Li, Q. Yang, and X. Huang, "A mobility-aware conflict-graph scheduling framework for auv-assisted collaborative underwater acoustic networks," *IEEE Internet of Things Journal*, pp. 1–1, 2026.

[10] R. Zhu, A. Boukerche, L. Feng, and Q. Yang, "A trust management-based secure routing protocol with auv-aided path repairing for underwater acoustic sensor networks," *Ad Hoc Networks*, vol. 149, p. 103212, 2023.

[11] E. P. M. Câmara Júnior, L. F. M. Vieira, and M. A. M. Vieira, "Scheduling nodes in underwater networks using voronoi diagram," in *Proceedings of the 20th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ser. MSWiM '17. Association for Computing Machinery, 2017, pp. 245–252.

[12] M. R. Jafri, A. Marin, A. Torsello, and M. Ghaderi, "On the optimality of opportunistic routing protocols for underwater sensor networks," in *Proceedings of the 21st ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, ser. MSWiM '18. Association for Computing Machinery, 2018, pp. 207–215.

[13] H. Niu, X. Li, Y. Zhang, and J. Xu, "Advances and applications of machine learning in underwater acoustics," *Intelligent Marine Technology and Systems*, vol. 1, no. 1, p. 8, 2023.

[14] H. Liu, L. Ma, Z. Wang, and G. Qiao, "Channel prediction for underwater acoustic communication: A review and performance evaluation of algorithms," *Remote Sensing*, vol. 16, no. 9, p. 1546, 2024.

[15] Y. Hu, S. Han, H. Li, H. Zhao, G. Yang, and J. Xu, "Tcn-based m-ary mobile spread spectrum underwater acoustic communication," *Applied Acoustics*, vol. 211, p. 109457, 2023.

[16] Y. Zhao, Y. Chen, W. Zhang, X. Wang, Y. Tao, C. Li, and X. Xu, "A multiagent reinforcement learning-based mac protocol for clustered multihop underwater acoustic sensor networks," *IEEE Sensors Journal*, vol. 25, no. 15, pp. 30 034–30 046, 2025.

[17] Y. Ma, C.-X. Wang, H. Chang, J. Huang, J. Wang, L. An, and Y. Chen, "A novel nonstationary geometry-based stochastic model for underwater acoustic mimo communication systems in shallow seas," *IEEE Internet of Things Journal*, vol. 12, no. 20, pp. 41 713–41 726, 2025.

[18] B. Zhao, J. Wu, Y. Ma, and C. Yang, "Meta-learning for wireless communications: A survey and a comparison to gnns," *IEEE Open Journal of the Communications Society*, vol. 5, pp. 1987–2015, 2024.

[19] A. Tziouvaras, B. Bertalanic, G. Floros, K. Kolomvatsos, P. Sarigiannidis, and C. Fortuna, "A representation learning approach to feature drift detection in wireless networks," *IEEE Transactions on Emerging Topics in Computing*, vol. 14, no. 1, pp. 80–93, 2026.

[20] E. S. Erstorp, V. Lidström, and P. Sigray, "Dlink: Introducing a framework for link adaptation in flooding-based underwater networks," *IEEE Journal of Oceanic Engineering*, vol. 50, no. 2, pp. 1456–1468, 2025.

[21] L. Feng, J. Liu, Y. Chen, and J. Yao, "A fair mac protocol based on dual bandwidth allocation iterations in underwater acoustic sensor networks," *International Journal of Communication Systems*, vol. 38, no. 4, p. e6017, 2025.

[22] J. Jiang, Y. Dong, G. Han, and G. Su, "Underwater acoustic mac protocol for multi-objective optimization based on multi-agent reinforcement learning," *Drones*, vol. 9, no. 2, p. 123, 2025.

[23] J. Huang, X. Ye, Y. Wang, and L. Fu, "Leveraging propagation delays: A delay-aware multi-agent reinforcement learning mac protocol for underwater acoustic networks," *IEEE Internet of Things Journal*, 2025.

[24] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *International conference on machine learning*. PMLR, 2017, pp. 1126–1135.

[25] X. Ye, L. Fu, X. Song, and Y. Wu, "Energy-efficient link adaptation for underwater acoustic communications based on meta deep reinforcement learning," *IEEE Internet of Things Journal*, 2025.

[26] P. van Walree, F.-X. Socheleau, R. W. Otnes, and T. Jenserud, "The watermark benchmark for underwater acoustic modulation schemes," *IEEE Journal of Oceanic Engineering*, vol. 42, no. 4, pp. 1007–1018, 2017.