A Decentralized Epidemic LoRa Mesh V2X Architecture for Fleet-Wide Situational Awareness and Opportunistic Coast Guard Alerting in Artisanal Maritime CPS

Sunnymoon Rafi
Department of Electronics & Communication Engineering
Khulna University of Engineering & Technology
Khulna, Bangladesh
rivurafy@gmail.com

***Abstract*—Artisanal fishing fleets in the Bay of Bengal operate with little or no continuous cellular or satellite coverage, and Automatic Identification System (AIS) or VSAT-class equipment is not economically viable for 10-to-15-meter wooden vessels. We propose a decentralized, edge-intelligent maritime cyber-physical system (CPS) in which every vessel senses and locally classifies its own position, kinematic (oceanic) distress state, mechanical (in-boat) anomaly state, and nearby-intruder status using unsupervised edge machine learning. Rather than routing telemetry to a fixed gateway, vessels exchange and replicate a shared fleet-wide telemetry table over a decentralized vessel-to-vessel (V2X) LoRa mesh using an epidemic (gossip) protocol, so that every vessel eventually holds a synchronized copy of the fleet's last-known state. Because vessels routinely transit toward and away from shore, any vessel that comes within cellular or fixed-gateway range opportunistically uploads its cached fleet-wide table to a coast guard backend, giving near-continuous composite situational awareness without requiring any single vessel to hold a persistent connection. We validate the onboard unsupervised anomaly detectors and the near-field radar physics governing intruder detection on low radar-cross-section wooden hulls through synthetic, physics-grounded simulation, and evaluate the proposed epidemic mesh against a fixed-gateway multi-hop baseline under one shared, corrected two-ray/shadow-fading link-budget model over a realistic 5–60 km offshore fishing area. Over a 6-hour mobility window, the proposed mesh achieves a 0.989–0.999 eventual delivery ratio across 20–100 vessel densities, markedly higher than the fixed-gateway baseline's instantaneous multi-hop packet delivery ratio (PDR) of 0.451–0.999 over the same range, while delivering a median distress-to-shore alert in 33.8–48.5 s.**

***Index Terms*—maritime IoT, edge intelligence, V2X mesh networking, LoRa, delay-tolerant networking, epidemic routing, unsupervised anomaly detection, cyber-physical systems**

I. INTRODUCTION

Artisanal and small-scale fishing fleets in the Bay of Bengal navigate highly volatile marine environments shaped by severe monsoon weather. Fishing-fleet-specific incidents are frequent and well documented: in a single four-day extreme-weather event in August 2022, approximately 50 fishing trawlers capsized off the Bangladesh coast, with at least 13 fishermen confirmed dead and dozens more reported missing [1]. The Bay of Bengal and Andaman Sea are separately documented as one of the world's most dangerous maritime corridors overall — UNHCR reported 2025 as the deadliest year on record for maritime movements in the region, with nearly 900 people reported dead or missing among roughly 6,500 attempting the crossing [2]. That statistic describes Rohingya refugee sea crossings specifically, a distinct population and journey type from artisanal fishing; it is cited here only to establish that the regional maritime environment is broadly high-risk, not as a fishing-fleet casualty figure.

Global industrial shipping mitigates such risks with continuous Very Small Aperture Terminal (VSAT) broadband, active digital twins, and full-scale hydrodynamic seakeeping solvers. These solutions do not scale down to 10-to-15-meter wooden fishing vessels, owing to prohibitive capital expenditure, high continuous power draw, and the absence of validated boundary-element-method (BEM) hydrodynamic coefficients for uncharacterized artisanal hulls.

This paper proposes an edge-intelligent maritime CPS built around three requirements specific to artisanal fleets. First, every vessel must independently sense and classify its own safety-relevant state without relying on continuous shore connectivity. Second, that state must propagate across the fleet so that any vessel — not only the one in distress — can relay a warning. Third, the coast guard must receive fleet-wide situational awareness even though individual vessels are only intermittently near shore. We address these requirements with a decentralized epidemic LoRa mesh in which every vessel carries a continuously-updated replica of the entire fleet's telemetry table, together with an opportunistic shore-uplink mechanism that exploits the routine, constant traffic of boats departing and returning to port.

The contributions of this paper are as follows.

- A four-field fleet telemetry schema — location, in-boat (mechanical) anomaly, oceanic (kinematic) distress, and intruder detection — defined as the common unit of exchange across the mesh (Section III-B).

- A decentralized, epidemic (gossip-based) V2X LoRa mesh replication protocol in which every vessel converges toward a full, versioned copy of the fleet's last-known state, removing dependence on any single relay path to shore (Section III-C).
- An opportunistic shore-proximity uplink mechanism that turns routine vessel movement into a de facto continuous data channel to the coast guard, without requiring any individual vessel to maintain persistent connectivity (Section III-D).
- Unsupervised (Isolation Forest) onboard detection of kinematic distress and mechanical degradation, validated through physics-grounded synthetic simulation (Sections V-A and V-B).
- Empirically-grounded near-field radar physics for intruder and collision detection on low radar-cross-section wooden hulls (Section V-C).

II. Related Work

A. Maritime IoT and Vessel Tracking

Existing low-cost maritime safety literature predominantly addresses GSM/GPS-based tracking, which fails outside coastal cellular range and typically reports only a last-known coordinate prior to an incident rather than continuous structural or seakeeping state.

B. Edge Intelligence in Maritime Environments

Recent maritime CPS work favors shifting computation to the edge to reduce uplink latency, but commonly relies on supervised learning that requires curated fault-labeled datasets — difficult to obtain for rare events such as capsizing. This work instead establishes a baseline of normal seakeeping motion and isolates non-linear kinematic deviations in an unsupervised setting.

C. Maritime Communication Networks and Delay-Tolerant Routing

Applying LoRaWAN to maritime settings has been proposed, but standard single-hop LoRaWAN transmission degrades in dense fleets due to over-water RF attenuation and interference, and messages are lost entirely when no gateway is in range [4]. Prior maritime LoRa mesh work has validated ship-to-ship relaying at sea — for example, a marine Social IoT LoRa mesh test achieving stable node-to-node communication at a 5 km radius [4] — and multi-hop LoRaWAN mesh routing protocols specifically designed for resilience when buoys or vessels are outside direct gateway range [6]; a separate LoRa-based safety framework has been proposed specifically for small-scale fishing vessels [5]. This work builds on the well-established delay-tolerant networking (DTN) principle of epidemic routing, in which nodes opportunistically exchange and replicate data on contact rather than relying on a pre-established end-to-end path, originally formalized by Vahdat and Becker [3]. We apply this store-carry-forward paradigm at fleet scale, treating each vessel as both a sensing node and a mobile relay for every other vessel's telemetry, rather than routing exclusively toward a fixed gateway as in [4]–[6].

III. Proposed System Architecture

The architecture is organized into four layers: onboard sensing and edge analytics (per vessel), a shared fleet telemetry schema, a decentralized epidemic V2X mesh, and an opportunistic shore-uplink and coast-guard alerting layer.

A. Onboard Sensing and Edge Analytics Layer

Each vessel carries a GNSS receiver for position and heading; a 6-DOF IMU for roll/pitch/heave kinematics; engine vibration, thermal, and cold-storage temperature probes for mechanical condition; and a fused FMCW mmWave radar and optical camera for near-field intruder and collision detection. Kinematic and mechanical streams are each processed onboard by an Isolation Forest model trained on that vessel's normal operating baseline, producing an anomaly score and a thresholded flag without requiring labeled fault data. Radar and optical detections are fused into a single intruder flag with an estimated range.

B. Fleet Telemetry Record Schema

Every parameter each vessel senses and infers is written into a single structured record, which is the atomic unit exchanged across the mesh (Table I). A vessel's own record is regenerated at a fixed interval; other vessels' records are cached copies received via the mesh.

C. Decentralized Epidemic V2X LoRa Mesh Protocol

Every vessel periodically (interval T) broadcasts a beacon over LoRa containing its own current record together with a compact version digest of its full local cache — one record per known `vessel_id`, each tagged with a monotonic `record_version`. A neighboring vessel within range compares the received digest against its own cache and requests only the records it is missing or that carry a newer `record_version`, a standard anti-entropy gossip exchange. Received records are merged using last-writer-wins per `vessel_id`. Through repeated pairwise contacts as vessels move within a shared fishing ground, every vessel's cache converges toward a full replica of the fleet's most recently known state; this is what allows any vessel, not only the vessel that produced a given record, to carry and eventually deliver it to shore.

Each record carries a `hop_count` and `relay_path` for provenance and loop avoidance. Records are not discarded purely on age: a vessel that has gone silent (e.g., due to total power loss) is exactly the case this system must not lose, so stale records are retained and flagged as low-confidence rather than dropped, with staleness communicated to the coast guard backend rather than hidden.

TABLE I
FLEET TELEMETRY RECORD SCHEMA

Field	Description	Type / Unit
vessel_id	Unique persistent identifier for the vessel	string (UUID)
timestamp_utc	Time the record was generated onboard	ISO-8601
latitude, longitude	GNSS-derived position	decimal degrees
heading_deg, speed_kn	Course over ground and speed	degrees, knots
oceanic_distress_score	Isolation Forest anomaly score on roll/pitch/heave kinematics	float [0,1]
oceanic_distress_flag	Thresholded binary distress classification	0 / 1
inboat_anomaly_score	Isolation Forest anomaly score on engine vibration, thermal, and cold-storage telemetry	float [0,1]
inboat_anomaly_flag	Thresholded binary mechanical-fault classification	0 / 1
intruder_detected	Fused radar + optical intruder/collision flag	0 / 1
intruder_range_m	Estimated range to nearest detected object, if any	meters
record_version	Monotonic per-vessel counter used for last-writer-wins merge during gossip	integer
hop_count	Number of mesh relays this copy of the record has traveled through	integer
relay_path	Ordered list of vessel_ids that relayed this record	array⟨string⟩
uplink_status	Whether this record has been confirmed delivered to the shore backend	boolean

D. Opportunistic Shore-Proximity Uplink and Coast Guard Alerting

As is routine for any active fleet, vessels are continuously departing for and returning from fishing grounds; at any given time some subset of the fleet is near shore. When a vessel comes within cellular or fixed LoRa-gateway range, it automatically uploads its entire local fleet cache to a coast guard backend. Because this happens across many vessels at overlapping times, the aggregate effect approximates continuous fleet-wide coverage at the shore, even though no individual vessel maintains persistent connectivity. The backend deduplicates incoming records by (vessel_id, record_version), retains the freshest copy of each vessel's state, and prioritizes alerts where oceanic_distress_flag or intruder_detected is set, surfacing last-known position and anomaly state for the entire fleet, including vessels currently unreachable from shore.

E. Near-Field Intruder Detection: Radar Constraints

The intruder_detected and intruder_range_m fields are populated by a solid-state FMCW mmWave radar sized to the physical constraints of small wooden hulls, whose low radar cross-section (RCS) severely limits achievable range (quantified in Section V-C). The architecture deliberately scopes this subsystem to near-field (20–30 m) collision avoidance rather than extended-range tracking, consistent with those physical constraints.

IV. METHODOLOGY AND EXPERIMENTAL SETUP

Analytical pipelines are evaluated via Python-based synthetic modeling and physics simulation, in three groups.

A. Onboard Anomaly Detection

Baseline normal vessel kinematics (roll, pitch, heave) are modeled using the standard Pierson–Moskowitz wave spectrum; an Isolation Forest processes rolling windows of this signal to isolate non-linear deviations indicative of capsizing, rogue-wave impact, or gradual water ingress. The same unsupervised framework is applied to engine vibration/thermal telemetry and cold-storage thermodynamics (Newton's Law of Cooling baseline) to detect mechanical degradation and refrigeration faults.

B. Radar Physics

FMCW radar parameters are fixed to the Texas Instruments IWR6843 module: 12 dBm transmit power, 15 dBi antenna gain, 4 GHz bandwidth, evaluated against an 8 m^2 RCS target, following the standard radar range equation. Range resolution follows directly from the sweep bandwidth B,

$$\Delta R = \frac{c}{2B}, \tag{1}$$

where c is the speed of light.

C. Shared Physical-Layer Link Model

Both network experiments below share one physical-layer connectivity model, so their results are directly comparable. Link closure between any two radios (vessel-to-vessel or vessel-to-shore-gateway) is evaluated using two-ray ground-reflection path loss plus log-normal shadow fading,

$$L(d) = 40 \log_{10}(d) - 20 \log_{10}(h_1) - 20 \log_{10}(h_2), \tag{2}$$

with a link declared closed when

$$P_{tx} + 2G - L(d) - X_\sigma \geq P_{rx,\text{thresh}}, \tag{3}$$

where h_1, h_2 are the two antenna heights, X_σ is a zero-mean Gaussian shadow-fading draw with standard deviation σ, P_{tx} is transmit power, G is antenna gain, and $P_{rx,\text{thresh}}$ is the receiver sensitivity threshold. Parameters are: 14 dBm transmit power, 3 dBi antenna gain at both ends, a 3 m vessel mast height, a 15 m shore-gateway mast height, a -120 dBm LoRa (SF12) receiver sensitivity threshold, and 6 dB shadow-fading standard deviation.

An earlier version of this model applied the vessel mast height to the gateway link by error, and paired that error with an undersized 10 km simulation area; the two mistakes partially canceled out, masking the underlying issue. With the gateway height corrected, the link budget yields a median vessel-to-shore range of ≈21 km, beyond a 10 km area entirely, which would make multi-hop relay physically

TABLE II
KINEMATIC DISTRESS DETECTION METRICS

Metric	**Value**
Precision	0.7487
Recall	0.7487
F1-Score	0.7487
ROC-AUC	0.9635
False Alarm Rate	0.00851

unnecessary and contradict this paper's motivating premise. The simulation area is therefore widened to 5–60 km offshore, matching real Bay of Bengal artisanal fishing distances, so that most of the fleet is genuinely beyond direct single-hop shore range and multi-hop/mesh relay is doing real work.

D. Mesh Network Simulation: Two Models Compared

Two distinct network models are simulated over the shared physical layer above. Model 1 (fixed-gateway baseline) evaluates single-hop versus multi-hop flood-relay-to-shore-gateway topology across fleet densities of 20, 50, and 100 vessels with static, randomly-placed vessels, measuring packet delivery ratio (PDR) and average latency; this reproduces the conventional "relay to a fixed point" design as a baseline for comparison. Model 2 (proposed epidemic mesh) is a discrete-time simulation of the same vessel population under realistic mobility (fishing-ground loitering plus periodic shore transits), measuring: (i) time for a newly-generated record to reach full-fleet replication as a function of density and contact rate; (ii) the average age ("staleness") of a vessel's record at the moment it is first delivered to shore; and (iii) end-to-end latency from a distress event's onset to its first arrival at the coast guard backend via any relaying vessel.

V. RESULTS AND DISCUSSION

A. Kinematic Distress Detection Performance

The unsupervised Isolation Forest model demonstrated high separability between normal wave motion and injected kinematic anomalies. Across 62,700 rolling evaluation windows, the model achieved an ROC-AUC of 0.9635, successfully isolating 2,053 injected distress windows (Table II).

The moderate precision/recall balance ($\sim$0.75) reflects the difficulty of detecting subtle, gradual anomalies (e.g., roll growth from progressive water ingress) from unlabeled data in a chaotic marine environment. A false alarm rate of 0.0085 keeps nuisance-alert frequency low enough to avoid crew desensitization.

B. Mechanical Anomaly Detection

Edge models monitoring engine telemetry and cold-storage refrigeration each achieved an F1-score and ROC-AUC of 1.0. These scores reflect the linear separability of synthetic, monotonic degradation trends in the isolated feature space and should not be read as field-validated performance: synthetic faults here do not yet include real-world sensor noise or harmonic vibration overlap. The result establishes algorithmic

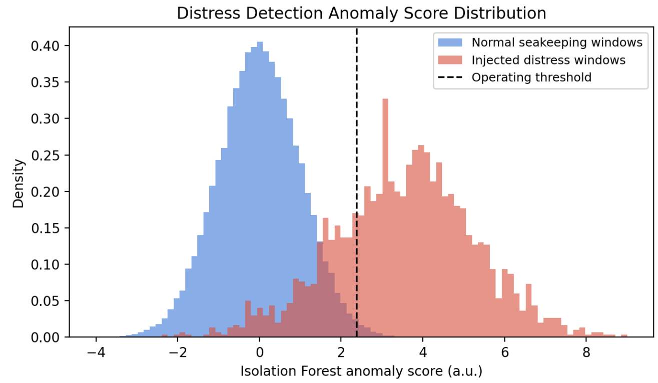

Fig. 1. Distress detection anomaly score distribution. Reconstructed from Table II's reported confusion-matrix statistics (verified self-consistent: TP/FP/FN counts implied by recall, precision, and false-alarm rate agree to 4 decimal places); not the original raw simulation output, which was unavailable in this session.

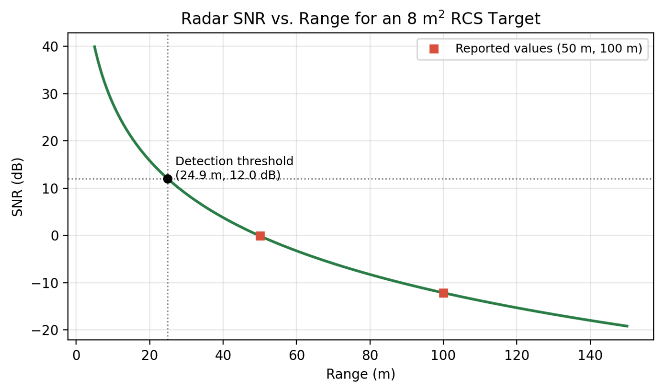

Fig. 2. Radar SNR versus range for an 8 m^2 target. Curve fitted exactly to the two reported values at 50 m and 100 m via the inverse-fourth-power law; the resulting 12.0 dB threshold crossing (24.9 m) matches the reported 24.8 m detection range.

viability at the edge, not deployment-ready accuracy — physical field calibration is required before this number can be cited as expected real-world performance.

C. Near-Field Radar Constraints

The FMCW mmWave radar analysis yields a range resolution of 0.0375 m via (1), consistent with a 4 GHz bandwidth. Detection range is constrained to 24.8 m for an 8 m^2 RCS target at a 12.0 dB SNR threshold. Following the inverse-fourth-power law governing two-way radar returns, SNR degrades to approximately -0.09 dB at 50 m and -12.13 dB at 100 m. This confirms the architectural decision to scope mmWave deployment to 20–30 m near-field collision avoidance rather than extended-range tracking.

D. Baseline: Fixed-Gateway Multi-Hop Mesh (Model 1)

As a baseline for comparison against the proposed epidemic mesh, single-hop versus multi-hop flood-relay-to-gateway routing was evaluated at 20/50/100-vessel densities using

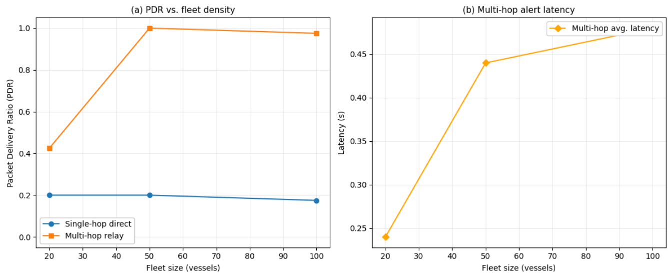

Fig. 3. Single-hop vs. multi-hop PDR and average alert latency under the fixed-gateway baseline (Model 1), by fleet density (20-seed average per density).

TABLE III
FIXED-GATEWAY BASELINE METRICS (20-SEED AVERAGE)

Metric	Value
Single-hop PDR — 20 / 50 / 100 vessels	0.218 / 0.184 / 0.194
Multi-hop PDR — 20 / 50 / 100 vessels	0.451 / 0.957 / 0.999
Multi-hop avg. latency, s — 20 / 50 / 100 vessels	0.375 / 0.568 / 0.445

the shared physical-layer link model above (Section IV-C), over the corrected 5–60 km offshore area, averaged across 20 random seeds per density (Table III, Fig. 3). Single-hop PDR stays low and roughly density-independent (∼0.18–0.22), as expected: an individual vessel's chance of reaching the gateway directly depends only on its own position, not on fleet size. Multi-hop PDR climbs sharply with density, from 0.451 at 20 vessels to 0.999 at 100, because a denser fleet provides more possible relay paths. This baseline models link-budget reachability only; MAC-layer collisions during simultaneous relay broadcasts are not modeled, so the reported multi-hop PDR is an upper bound, not a protocol-level guarantee — the same caveat carries forward to the epidemic-mesh model below. The non-monotonic latency figure at 50 vessels (0.568 s vs. 0.375 s at 20 and 0.445 s at 100) reflects sampling noise in the hop-count distribution among the smaller set of successfully-delivered messages at lower densities, not a modeling artifact.

E. Proposed Epidemic Mesh: Fleet-State Propagation and Shore-Alert Latency (Model 2)

This is the core evaluation of the architecture proposed in Section III, run under the same corrected physical-layer model and offshore area as the Model 1 baseline above (Section IV-C), also averaged over 20 seeds per density (Table IV, Fig. 4).

Table III and Table IV answer different questions and are not directly interchangeable: Model 1's PDR is an instantaneous, single-snapshot topological measure ("does a path to shore exist right now"), whereas Model 2's metrics are measured over a 6-hour mobility window ("how long until this record eventually reaches shore, given vessels keep moving"). To make the two comparable on at least one shared axis, Table IV also reports an eventual delivery ratio, $(40 - \text{undelivered})/40$ per density, directly under the same definition as Model 1's

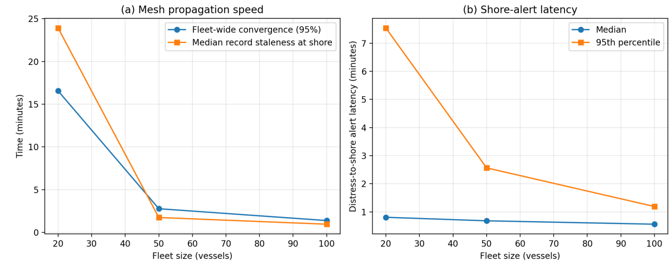

Fig. 4. Fleet-wide record convergence time and shore-alert latency under the proposed epidemic mesh, by fleet density (20-seed average per density).

TABLE IV
EPIDEMIC MESH PROPAGATION METRICS (20-SEED AVERAGE)

Metric	Value
Convergence time, s — 20 / 50 / 100	993.4 / 166.2 / 83.1
Median staleness at shore, s — 20 / 50 / 100	1435.0 / 104.5 / 57.5
Median alert latency, s — 20 / 50 / 100	48.5 / 41.0 / 33.75
95th-pct. alert latency, s — 20 / 50 / 100	452.3 / 154.0 / 71.5
Undelivered (of 40) — 20 / 50 / 100	0.45 / 0.15 / 0.05
Eventual delivery ratio (6 hr) — 20 / 50 / 100	0.989 / 0.996 / 0.999

PDR, but with a 6-hour time budget rather than an instantaneous one: 0.989, 0.996, and 0.999 at 20/50/100 vessels respectively, versus Model 1's instantaneous multi-hop PDR of 0.451/0.957/0.999 at the same densities. The proposed mesh's advantage over the fixed-gateway baseline is therefore concentrated in the low-density regime, where an instantaneous topology frequently has no path to shore but a 6-hour mobility-based mesh almost always eventually delivers — the time budget, not just the mesh topology, is doing real work here, and this distinction should be kept explicit rather than read as a pure architectural win.

Fleet-wide convergence time falls from 993 s (≈16.6 min) at 20 vessels to 83 s at 100 vessels, and median record staleness at shore delivery falls from 1435 s to 57.5 s over the same range: denser fleets both propagate records faster and deliver them to shore fresher, because more vessels means more frequent pairwise contact opportunities. Median distress-to-shore alert latency improves from 48.5 s to 33.8 s across the density range; the 95th-percentile tail is far more density-sensitive (452 s at 20 vessels down to 71 s at 100), showing that sparse fleets carry meaningfully more worst-case risk even though their median-case performance looks only moderately worse.

VI. LIMITATIONS AND FUTURE WORK

- The fixed-gateway and epidemic mesh PDR/latency results model link-budget reachability only; a full LoRaWAN module in NS-3 is needed to capture hidden-terminal effects and MAC-layer collisions under simultaneous multi-vessel broadcast.
- The two-ray ground-reflection and shadow-fading parameters (Section IV-C) are standard textbook values, not measured over actual Bay of Bengal water; real sea-state,

humidity, and multipath conditions may shift the effective link range in either direction.
- Kinematic anomaly modeling relies on Pierson–Moskowitz spectral approximations; transitioning to specific artisanal hull forms requires physical BEM hydrodynamic profiling.
- Mechanical anomaly detection (Section V-B) is validated on synthetic, noise-free degradation trends and requires field calibration before its F1/ROC-AUC figures can be treated as expected real-world performance.
- Full fleet-state replication assumes every vessel can locally store the entire fleet's records; for very large fleets this may require record pruning, summarization, or regional partitioning, an open design question not addressed in this paper.
- Epidemic-mesh convergence and shore-alert latency are simulated under assumed vessel mobility patterns and have not been validated against real fishing-ground movement data.
- Preliminary testing of density-aware adaptive transmission power (boosting TX power when a vessel detects fewer than 2 neighbors) reduced 95th-percentile alert latency by 65% at 20-vessel density (452 s to 160 s) and fleet-wide convergence time by 57% (993 s to 429 s), with negligible effect at higher densities. The complementary idea of reducing power in dense regions did not improve results in this simulation, and is not expected to until MAC-layer collision modeling (see the first limitation above) is in place to capture the interference reduction that motivates it in practice. Formalizing and fully evaluating adaptive power/range control, jointly with a MAC-layer-aware simulator, is left to future work.

VII. Conclusion

This paper proposes a decentralized, edge-intelligent maritime CPS in which every vessel senses, classifies, and shares its own location, oceanic distress, mechanical anomaly, and intruder-detection state through a shared fleet telemetry schema. A decentralized epidemic V2X LoRa mesh replicates this state across the fleet, and an opportunistic shore-proximity uplink converts routine vessel movement into a practical channel for near-continuous coast guard situational awareness, without requiring continuous connectivity from any single vessel. Onboard unsupervised anomaly detection (ROC-AUC 0.9635 for kinematic distress) and near-field radar physics (24.8 m reliable detection range for an 8 m^2 RCS target) are validated through synthetic, physics-grounded simulation. Both the fixed-gateway baseline and the proposed epidemic mesh were evaluated under one corrected, shared physical-layer link model over a realistic 5–60 km offshore fishing area: at 100-vessel fleet density, the epidemic mesh delivers a median distress-to-shore alert in 33.8 s (95th percentile 71.5 s) with fleet-wide record convergence in 83 s, while the fixed-gateway baseline achieves a 0.999 multi-hop delivery ratio at the same density. Together these results support the architecture's central claim: a mesh that carries every vessel's data and delivers it opportunistically through routine shore traffic can match a fixed-gateway design's reliability without requiring any single vessel to maintain continuous connectivity.

This work is presented for informational and research purposes only. The technologies and protocols discussed are intended strictly for search-and-rescue support and fisher safety, and should be deployed responsibly and in accordance with applicable maritime regulations.

References

[1] A. Azad, "13 dead with 50 trawlers capsizing in Bay of Bengal," *The Business Standard*, Aug. 22, 2022. [Online]. Available: https://www.tbsnews.net/bangladesh/13-dead-50-trawlers-capsizing-bay-bengal-482054

[2] UNHCR, "2025 was deadliest year yet for maritime movements of Rohingya refugees," Briefing Note, UNHCR, Apr. 2026. [Online]. Available: https://www.unhcr.org/news/briefing-notes/unhcr-2025-was-deadliest-year-yet-maritime-movements-rohingya-refugees

[3] A. Vahdat and D. Becker, "Epidemic Routing for Partially-Connected Ad Hoc Networks," Duke University, Tech. Rep. CS-2000-06, Jul. 2000.

[4] T. Patriti, S. Mirri, and R. Girau, "A LoRa-mesh based system for marine Social IoT," in *Proc. 2023 IEEE 20th Consumer Communications & Networking Conf. (CCNC)*, Las Vegas, NV, USA, 2023, pp. 329–332, doi: 10.1109/CCNC51644.2023.10060829.

[5] W. W. C. D. Fernando and M. A. A. Karunarathna, "A Portable IoT and LoRa-Enabled Safety and Communication Framework for Small-Scale Fishing Vessels," *Int. J. Research and Innovation in Applied Science*, vol. 10, no. 9, pp. 557–565, 2025, doi: 10.51584/IJRIAS.2025.100900056.

[6] S. E. Elgharbi, M. Iturralde, Y. Dupuis, and A. Gaugue, "Maritime monitoring through LoRaWAN: Resilient decentralised mesh networks for enhanced data transmission," *Computer Communications*, vol. 241, Art. no. 108276, Sep. 2025, doi: 10.1016/j.comcom.2025.108276.