Power-Efficient LEO Satellite AIS Demodulator

Aditya Varma Gadiraju, Jyothendar Paladugula and Sriram Jayasimha
Signion Systems
Hyderabad, India
aditya@signion.com, jyothendar@signion.com and sriram@signion.com

***Abstract*—On-board satellite receivers for an Automatic Identification System (S-AIS) use a parallel demodulator structure (PDS) that executes a Doppler-timing search for every reception window, resulting in a computational cost that is independent of maritime traffic density. Much of a low-Earth orbit (LEO) satellite's ground track lies over sparse ocean where the PDS is under-utilized; the receiver could save on-board energy by selecting adaptively between low-count instantiations of per-burst demodulators and the reference PDS bank at each two-second reception window using two block-level metrics computed before demodulator execution. These are the ratio of block RMS to a Rayleigh noise estimate, and the count of preamble matched filter peaks detected above the noise threshold. Performance is evaluated using two physics-based satellite AIS scenarios that are AI generated with verification to comply with ITU-R M.1371 and M.2169. The adaptive scheduler recovers 67 of 67 and 182 of 184 class-A vessels detected by the reference receiver while reducing computational effort by 76.3% and 27.8%, respectively. The results show that computational effort, and therefore on-board receiver power consumption, is reduced for sparse regions with the same effective per-pass recovery as the reference receiver.**

Keywords—Automatic Identification System, SAIS, satellite communications, on-board power efficiency, adaptive scheduling, GMSK, PDS, SIC.

I. INTRODUCTION

The Automatic Identification System (AIS) is a maritime self-organizing time-division multiple-access (SOTDMA) protocol operating in two or more Very High Frequency (VHF) channels that broadcasts vessel navigation data [1]. Space-based AIS on low-Earth-orbit (LEO) satellites extends surveillance globally, observing a field of view (FOV) up to a 2000 km radius with Doppler offsets up to ±4 kHz and varying arrival times [2].

Because SOTDMA only coordinates locally, overlapping cells cause frequent packet collisions at the satellite [3], [4]. Standard receivers use a parallel demodulator structure (PDS) [5]: a fixed bank of 21 Doppler by six timing-offset non-coherent GMSK demodulators. While robust, PDS incurs a high, constant computational cost regardless of traffic density.

Although vessel density drops substantially over open oceans, sparse satellite-AIS regions cannot be treated as conventional terrestrial scenarios. The wide FOV aggregates geographically separated SOTDMA cells, producing collisions even at low densities. Simply duty-cycling over sparse regions reduces computation but risks missing critical non-cooperative or non-compliant activity.

Existing work addresses dense-collision processing [5], [6], [7], [8], ground-station offloading [9], or mission-level duty-cycling based on orbital constraints. However, block-level adaptation of receiver complexity to maritime traffic density remains unaddressed. This motivates adaptive processing that preserves full observational coverage while redirecting saved energy to downlink, station-keeping, or additional satellite payload operations.

An alternative architecture is the per-burst (PB) demodulator: a preamble matched filter localizes candidate transmissions in the delay-Doppler plane, applying demodulation only near detected peaks. In sparse scenes, PB drastically cuts computation compared to PDS. However, PB is fundamentally peak-driven; in dense scenes with overlapping bursts, numerous signal-like peaks cause PB complexity to approach or exceed PDS.

We introduce a novel adaptive scheduler that selects between PB, the full-bank PDS (REF), and SKIP for each 2-s reception window using two pre-demodulation metrics. Thresholds are derived from Rayleigh noise statistics and per-block cost parity between PB and REF. The contributions are: (i) a two-metric, physics-derived scheduler using the pre-computed block RMS ratio and strong preamble peak count; (ii) a physics-based synthetic evaluation scenario adhering to ITU-R M.1371 and M.2169; and (iii) characterization of the compute-recovery operating point identifying where PB outperforms REF.

II. RELATED WORK

A. Parallel demodulator structures

Burzigotti et al. [3] introduced the PDS for satellite AIS using non-coherent continuous-phase-modulation sequence detection [4]. Jayasimha et al. [5] extended the PDS with a multi-stage pipeline aimed at dense maritime environments. These receivers use a fixed PDS for every reception window that does not adapt to per-window signal statistics.

B. Preamble-based receivers

Gallardo and Sorger [10] described a coherent satellite-AIS receiver in which a preamble matched filter is used together with coherent demodulation. In such receivers, the preamble matched filter is used for message detection, not for scheduling between competing demodulator architectures.

C. Doppler and timing estimation

Meng et al. [11] used interpolated fast Fourier transform methods to satellite-AIS Doppler estimation. Morelli and Mengali [12] derived joint frequency and timing recovery for minimum-shift-keying type modulation. Rife and Vincent's classical parabolic (magnitude-squared) peak-interpolation method [13] is adopted in the present work to refine frequency and timing after matched-filter peak detection.

D. Multi-user and co-channel processing

For completeness, Picard et al. [6] proposed a multi-user, multi-antenna satellite-AIS receiver that exploits direction-of-arrival, polarization and Doppler diversity with a payload size penalty, while Nelson et al. [14] addressed coherent AIS demodulation in the presence of dense-zone co-channel interference, with performance obtained at the expense of parallel processing of many phase hypotheses. This in no way vitiates our proposed sparse-zone computational savings.

E. Successive Interference Cancellation

Successive interference cancellation (SIC) has been investigated to improve the recovery of colliding AIS bursts in dense maritime environments. Jayasimha et al. [5] incorporated an iterative SIC stage within the PDS, where successfully decoded bursts are reconstructed and subtracted from the composite received signal to enable recovery of weaker overlapping transmissions. Uehashi et al. [7] further extended SIC for asynchronous signal collisions in space-based AIS receivers, demonstrating improved burst recovery under realistic satellite reception conditions. Since the proposed adaptive scheduler determines receiver selection rather than the underlying burst demodulation process, SIC is complementary to the proposed architecture and can be integrated into both the PB and reference receiver paths to further improve recovery performance.

III. Signal Model and Physics-Based Scenario

A. Received signal

The complex baseband received signal $r[n]$ in one 2-second reception block is the superposition of K GMSK-modulated bursts with additive complex Gaussian noise:

$$r[n] = \sum_{k=1}^{K} \alpha_k s_k[n - \tau_k] \exp\left(\frac{j2\pi f_{d,k} n}{F_s} + j\varphi_k\right) + w[n] \quad (1)$$

for $n = 0,\ldots,N-1$ with $N = 115{,}200$ samples at $F_s = 57.6$ kHz. Here $s_k[n]$ is a GMSK-modulated burst with BT = 0.4, 0.5 modulation index, and 9600 bit/s bit rate [1]; α_k is the complex amplitude, τ_k the integer arrival time, $f_{d,k} \in [-4,+4]$ kHz the Doppler offset, φ_k an unknown carrier phase, and $w[n]$ is circularly symmetric complex Gaussian noise with variance σ^2. The channel-A carrier frequency is 161.975 MHz.

B. Link-budget model (ITU-R M.1371 / M.2169)

Physics-derived per-burst E_b/N_0 (dB) is:

$$E_b/N_0 = P_{tx} + G_{ant}(\theta_{zen}) - L_{impl} - L_{FSPL}(R) + K_{CAL} \quad (2)$$

where:

$$L_{\mathrm{FSPL}}(R) = 32.44 + 20\log_{10} R_{\mathrm{km}} + 20\log_{10} f_{\mathrm{MHz}}\ [dB] \quad (3)$$

$$G_{ant}(\theta_{zen}) = max(2.15 + 20\log_{10}\sin\theta_{zen}, -12)\ [dBi] \quad (4)$$

P_{tx} is the ship EIRP (10.97 dBW for Class A, 3.01 dBW for Class B [1]); G_{ant} is the $\lambda/2$ dipole antenna gain towards the satellite as a function of the zenith angle θ_{zen} [15]; L_{impl}=3 dB is the ship-side cable and implementation loss; L_{FSPL} is the free-space path loss at slant range R; K_{CAL} =145.84 dB is calibrated so that the physically achievable maximum E_b/N_0 (at $\theta_{\mathrm{el}} \approx 45°$) is approximately 17 dB, matching typical peak reception figure in Rec. ITU-R M.2169. The elevation angle at the ship for a satellite at altitude h and slant range R is:

$$\sin\theta_{el} = [(R_E + h)^2 - R_E^2 - R^2]/(2 \cdot R \cdot R_E) \quad (5)$$

with R_E= 6371 km and h = 638 km, and $\theta_{\mathrm{zen}} = 90° - \theta_{\mathrm{el}}$.

Table I lists anchor E_b/N_0 values given by this link budget.

C. AI based scenario generation

AIS scenarios are generated to evaluate the proposed adaptive scheduler under representative sparse and dense maritime operating conditions, using a two-stage AI-assisted pipeline. In the first stage, a vision-capable large language model [16] interprets publicly available satellite-AIS density maps [17] to extract the recurring spatial structure of maritime traffic, including coastal corridors, cross-ocean great-circle lanes, and open-ocean sparsity regions. From this interpretation, the model outputs a compact set of geographic descriptors consisting of latitude and longitude corridor bounding boxes, coastal offshore-decay parameters, and per-region vessel-density tiers that reproduce the observed density distribution. In the second stage, a physics-based simulator samples vessel positions from these AI-extracted regions and generates the corresponding AIS signal layer, including SOTDMA slot allocation with slot-contention statistics consistent with ITU-R M.1371, GMSK burst generation at 9600 bit/s, and SGP4-propagated satellite geometry and Doppler shift. The per-burst E_b/N_0 is evaluated from the ITU-R M.1371 and M.2169 link budget, accounting for vessel EIRP, the ship's dipole antenna gain in the satellite direction, and the satellite receiver G/T with the free-space path loss and slant range. The resulting channel-A and channel-B IQ records serve as input to the receiver analyzed in Section IV. This structure decouples where ships are placed, which is data-driven from real global maps, from how they transmit, which is physics-driven. Each resulting scenario is verified for specification compliance before use (by data binning vessel density per region, Class A/B distribution, vessel transmit intervals, and whether the computed values of EIRP and Doppler fall within the expected range), ensuring physical realism while substantially increasing scenario diversity and generation efficiency, enabling a comprehensive evaluation across a wide operating envelope.

TABLE I. Link-budget operating points at altitude 638km, matched to ITU-R M.2169 realistic reception figures.

Position	θ_{el}	R (km)	Class-A	Class-B	M.2169 target
Nadir (dipole null)	90°	638	9.0 dB	1.0 dB	A:8–10 B:0–2
Mid-FOV	27.7°	1200	16.7 dB	8.7 dB	A:16–20 B:8–12
Physical maximum	≈45°	~908	17.2 dB	9.2 dB	≤18 dB peak
Edge of FOV	10.3°	2000	13.2 dB	5.2 dB	(FOV-cap)

The simulation configuration is common to both scenarios. Vessel reporting interval follows the AIS SOTDMA protocol, with Class-A vessels operating in four operating modes (fast: 1.5-2.5 s, nominal: 5-7 s, slow: 9-11 s) and Class-B vessels transmitting every 28-32 s in accordance with the AIS specification [1]. Burst scheduling enforces slot uniqueness within 0.4° × 0.4° SOTDMA cells while preserving inter-cell time-frequency collisions observed from the satellite. Bursts are retained when the vessel is inside the 2000-km slant-range FOV of the satellite ephemeris.

Fig. 1 illustrates the South Pacific sparse scenario over a 40s observation interval. A total of 115 vessels are distributed uniformly by area within the satellite FOV, comprising 73% Class-A and 27% Class-B vessels. The resulting scenario contains 335 channel-A bursts, with received E_b/N_0 ranging from 4.3dB to 16.9dB and a median of 13.4dB, while Fig. 2 shows the Indian Ocean sparse-to-dense scenario, capturing a transition from open ocean to a busy shipping corridor during a 120-second pass. Over this pass, 204 class-A vessels remain within the satellite FOV and produce 2,965 channel-A bursts. This scenario is designed to stress the scheduler's mode-selection logic: the traffic density grows monotonically along the pass, so the density gauge is expected to select the per-burst path early in the pass and escalate to the reference bank in the dense corridor.

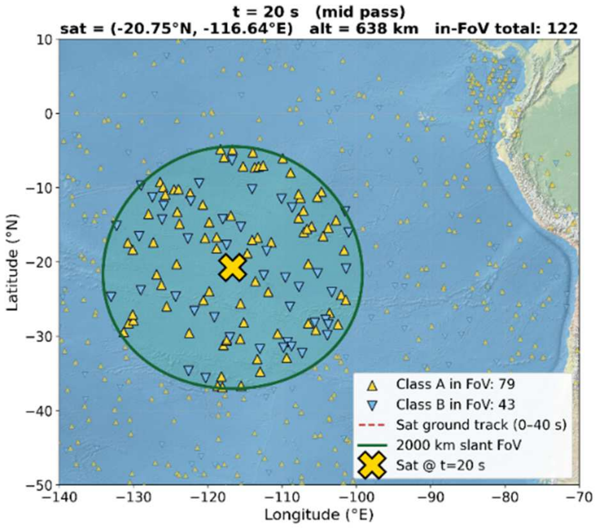

Fig. 1. South Pacific sparse scenario showing the satellite ground track and vessel distribution within the FOV (Class-A: yellow, Class-B: blue)

IV. Adaptive Receiver Architecture

Fig. 3 illustrates the proposed adaptive receiver architecture. Each 2s reception block is processed by a common front-end comprising signal pre-processing, preamble detection, and traffic-density estimation. Based on the estimated traffic density, the scheduler selects one of three receiver modes: SKIP, PB, or the reference PDS (REF).

A. Preamble matched filter

The received baseband signal is cross-correlated with 21 Doppler-shifted preamble templates:

$$p_m[n] = p[n]\exp(j2\pi f_d(m)n/F_s), \qquad (6)$$

$$f_d(m) = -4000 + 400 \cdot m \text{ Hz}, m = 0, \dots, 20$$

$$C[m,\tau] = \left|\sum_{n=0}^{L_p-1} r[\tau+n]p_m^*[n]\right| \qquad (7)$$

with L_p = 144 samples. Under the noise-only hypothesis, the matched filter magnitude $|C|$ follows a Rayleigh distribution with scale parameter σ_n proportional to the noise standard deviation. The Rayleigh assumption is well-supported for the satellite AIS channel: the footprint spans thousands of kilometres and observes many independent asynchronous transmissions, so by the Central Limit Theorem the envelope is Rayleigh-distributed [8], [18]. The peak-detection threshold is set from the 5th percentile of $|C|$:

$$\gamma_{th} = \gamma \cdot p_{5\%} \,/\, 0.3203, \gamma = 6.5 \qquad (8)$$

where $0.3203 = [-2 \cdot \ln(0.95)]^{0.5}$. Peaks are extracted iteratively: the maximum correlation peak of $|C|$ above γ_{th} is taken, then a guard region of ±1440 samples × ±1 Doppler bin around the peak is zeroed. Parabolic refinement across three cells nearest each peak is applied in both the frequency and time directions using the classical parabolic interpolation formulas of Rife and Vincent [13]:

$$\delta_f = \frac{|C_{-1}|^2 - |C_{+1}|^2}{2(|C_{-1}|^2 - 2|C_0|^2 + |C_{+1}|^2)}; \qquad (9)$$

$$\hat{f}_d = f_d + \delta_f \cdot 400 \text{ Hz}$$

$$\delta_\tau = \frac{|C(\tau-1)|^2 - |C(\tau+1)|^2}{2[|C(\tau-1)|^2 - 2|C(\tau)|^2 + |C(\tau+1)|^2]}; \qquad (10)$$

$$\delta_\tau \in [-0.5, +0.5] \text{ samples}$$

The refined estimates $(\bar{\tau}, \hat{f}_d)$ are computed by adding the parabolic corrections of (9)-(10) to the coarse cell values, and they replace the coarse cell-grid values in every downstream stage: $\hat{f}_d$ is used to de-rotate the burst slice to zero carrier prior to demodulation, and $\bar{\tau}$ sets the sample offset from which the 1440-sample slice is extracted. Each peak carries strength = $|C|/\gamma_{th} \geq 1$ by construction.

B. Density gauge — the scheduler decision

The scheduler operated on two block-level metrics computed prior to demodulation:

$$rms\_ratio = \frac{\sqrt{\text{mean}(|r|^2)}}{\text{median}(|r|)/1.1774} \qquad (11)$$

$$n_{\text{peaks}} = |\{\text{peak: strength(peak)} \geq 1\}| \qquad (12)$$

The factor $1.1774 = (\ln 4)^{0.5}$ converts a Rayleigh median to a Rayleigh scale parameter. Under the noise-only hypothesis, $rms\_ratio \to \sqrt{2} \approx 1.414$ and n_{peaks} reduces to the small residual count of γ_{th} false alarms.

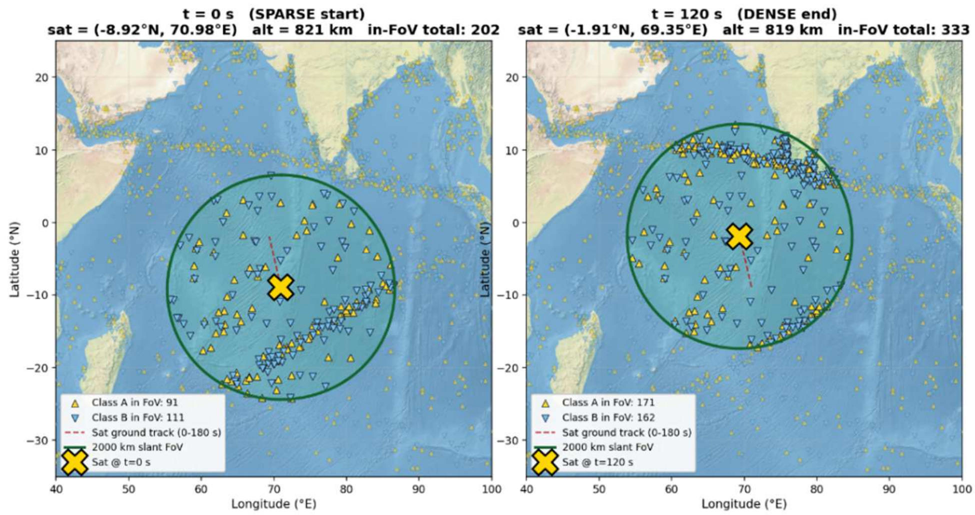

Fig. 2. Indian ocean sparse-to-dense satellite AIS scenario showing the transition from open-ocean traffic to a high-density shipping lane

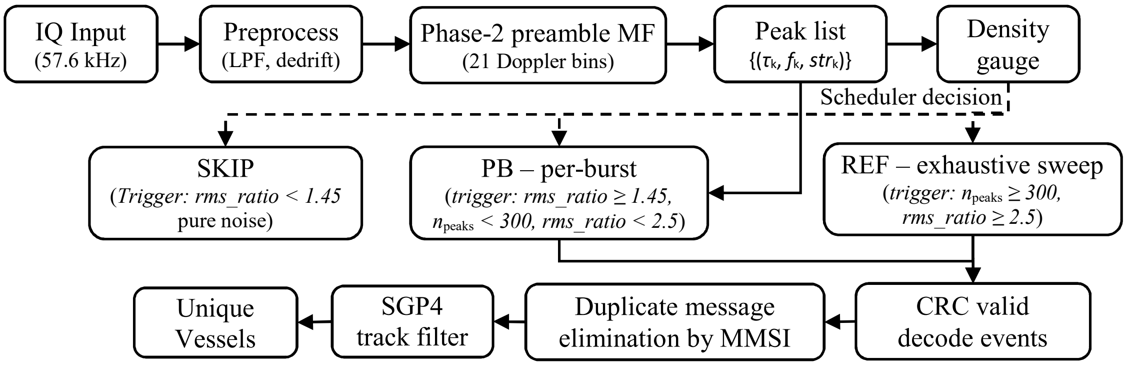

Fig. 3. Adaptive receiver architecture with scheduler selection between SKIP, PB, and REF processing paths

The scheduler decision is:

SKIP: *rms_ratio* < 1.45 (empty, noise-only)

PB: $n_{\text{peaks}} < 300$ and $1.45 \leq$ *rms_ratio* < 2.5 (sparse)

REF: $n_{\text{peaks}} \geq 300$ or *rms_ratio* ≥ 2.5 (dense)

The three thresholds are derived analytically. *rms_ratio* = 1.45 corresponds to 2.5% above the pure-noise limit √2 and rejects empty blocks. *rms_ratio* = 2.5 is a tuned threshold above which a significant fraction of the block contains burst energy. $n_{\text{peaks}} = 300$ is the compute-parity trigger derived in the section below.

C. Per-burst demodulator and reference bank

For each peak in a PB block, demodulation executes over a 2040-sample slice extracted around the refined peak location $(\bar{\tau}, \widehat{f_d})$. This slice includes a 200-sample lead for Viterbi state estimator convergence, a 1440-sample burst payload, and a 400-sample tail providing headroom for HDLC end-flag scanning. To absorb peak-location errors introduced by adjacent-burst interference, demodulation is repeated over a localized delay-Doppler sweep within this slice. The sweep covers 7 Doppler bins spaced by 400 Hz around $\widehat{f_d}$. The timing sweep τ_c adapts to the peak strength s: strong peaks ($s \geq 2.5$) maintain tight correlator lock and evaluate only $\tau_c \in \{0\}$ (7 calls). Weak peaks ($s < 2.5$) sweep $\tau_c \in \{-1000, 0, 1000\}$ samples (21 calls). Aggregated across the block, the total computational cost of the PB path in sample operations is:

$$C_{\text{PB}} = N_{\text{slice}} \cdot n_{\text{peaks}} \cdot (21 - 14 \cdot r_{2.5}) \quad (13)$$

where $N_{\text{slice}} = 2040$ is the per-slice sample count processed by the Viterbi and CRC stages, n_{peaks} is total peak count, and $r_{2.5} = |\{s \geq 2.5\}|/n_{\text{peaks}}$ is the block-level strong-peak fraction. Conversely, the REF path invokes demodulation across a fixed grid of 21 Doppler bins by 6 timing offsets over the entire block of length $N = 115{,}200$ samples. Its computational cost is constant regardless of traffic density:

$$C_{\text{REF}} = 21 \cdot 6 \cdot N = 14.5152 \times 10^6 \text{ sample ops} \quad (14)$$

HDLC framing allows REF to recover bursts anywhere in the block without requiring a preamble matched-filter peak, its core structural advantage in dense scenes and source of redundant computation in sparse ones. Setting $C_{\text{PB}} = (2/3) \cdot C_{\text{REF}}$ as the cost-parity threshold and assuming an empirical median strong-peak fraction of $r_{2.5} \approx 0.35$ yields $n_{\text{peaks}} \approx 300$, establishing the scenario-based compute-parity trigger for mode selection. Duplicate valid CRC decodes from all paths are eliminated by Maritime Mobile Service Identity (MMSI), while positional consistency is checked via a Simplified General Perturbations 4 (SGP4) satellite orbit propagator.

V. Results and Discussion

The proposed adaptive receiver was evaluated on both the physics-based synthetic scenarios described in Section III. For each scenario, all 2s blocks were processed sequentially by the adaptive scheduler and, for comparison, independently by the exhaustive reference bank (REF). Evaluation focuses on unique vessel recovery, total sample work, and per-class recovery rates. Table II summarizes the performance metrics across both the scenarios.

TABLE II. Performance Comparison of the Adaptive Scheduler and the Reference Receiver

Metric	South Pacific		Indian Ocean	
	Adaptive	**Reference**	**Adaptive**	**Reference**
Unique ships recovered	77	77	230	267
Class A recovery	67	67	182	184
Recovery ratio vs Reference	100%	100%	98.9%	100%
Total sample-work ($\times 10^6$)	68.8	290.3	628.8	870.9
Compute ratio (Adaptive / REF)	0.237	1.000	0.722	1.000
Cost per ship ($\times 10^6$)	0.89	3.77	2.73	3.26
Scheduler mode picks	20 PB + 0 REF	N/A	29 PB + 31 REF	N/A

A. Overall Performance

Figs. 4 and 5 illustrate the cumulative Class-A vessel recovery versus normalized computational effort for the South Pacific and Indian Ocean scenarios, respectively. In both cases, the adaptive scheduler closely tracks the reference recovery trajectory while significantly reducing computational work.

In the South Pacific sparse scenario, the scheduler continuously selects the low-complexity PB mode (20 of 20 blocks), achieving 100% Class-A vessel recovery relative to REF while reducing total computational effort by 76.3% (a compute ratio of 0.237). In the Indian Ocean transition scenario, traffic density rises along the pass. The scheduler dynamically adapts, selecting PB for 29 sparse/moderate blocks and escalating to REF for 31 dense corridor blocks. This yields a 27.8% overall compute reduction (compute ratio of 0.722) while recovering 182 of 184 Class-A vessels (98.9% relative recovery). These results confirm that the pre-demodulation density gauge scales complexity dynamically without sacrificing vessel recovery.

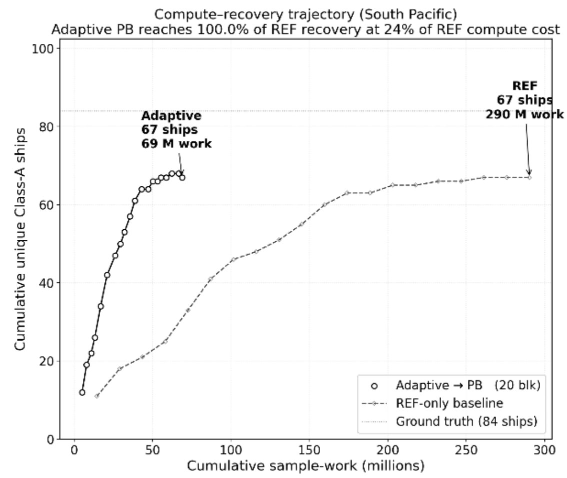

Fig. 4. Cumulative vessel recovery and normalized computational effort for the South Pacific sparse scenario

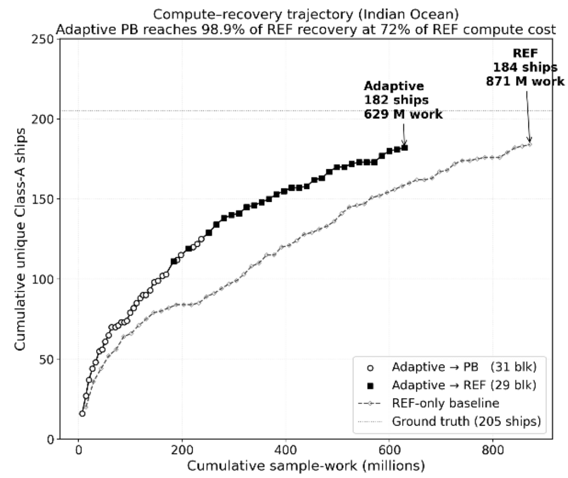

Fig. 5. Cumulative vessel recovery and normalized computational effort for the Indian Ocean sparse-dense scenario

B. Vessel recovery over a complete satellite pass

While individual AIS bursts may occasionally be lost to noise or co-channel interference, periodic retransmissions ensure high net vessel detection across an orbital pass. In the Indian Ocean scenario, 53 vessels remain within the satellite FOV for the full 2-minute duration; the adaptive scheduler successfully detects all 53 vessels (100% pass-level recovery). This confirms that burst-level dropouts in low-complexity modes do not degrade pass-level maritime surveillance reliability.

VI. CONCLUSION

This paper presented a novel adaptive demodulator scheduler for on-board satellite AIS receivers that dynamically selects between SKIP, per-burst, and REF processing using two physics-derived block-level metrics computed prior to demodulation. By adapting the receiver architecture to the observed traffic density, the proposed scheduler substantially reduces computational effort for full sparse-zone monitoring, while preserving decoding performance in dense traffic through automatic selection in the reference receiver.

Evaluation using two AI-assisted, physics-based satellite AIS scenarios compliant with ITU-R M.1371 and M.2169 demonstrated that the scheduler maintains recovery performance close to the reference receiver while reducing computational effort by up to 76.3% in sparse operating conditions and 27.8% during a sparse-to-dense satellite pass. These results demonstrate that adaptive scheduling improves computational efficiency without compromising vessel-level recovery, making it well suited for power-constrained on-board satellite receivers.

REFERENCES

[1] Recommendation ITU-R M.1371-5, "Technical characteristics for an automatic identification system using time-division multiple access in the VHF maritime mobile band," 2014.

[2] Report ITU-R M.2169, "Improved satellite detection of AIS," 2009.

[3] P. Burzigotti, A. Ginesi, and G. Colavolpe, "Advanced receiver design for satellite-based automatic identification system signal detection," Int. J. Satell. Commun. Netw., vol. 30, no. 2, pp. 52–63, 2012.

[4] G. Colavolpe and R. Raheli, "Noncoherent sequence detection of continuous phase modulations," IEEE Trans. Commun., vol. 47, no. 9, pp. 1303–1307, 1999.

[5] S. Jayasimha, J. Paladugula, A. V. Gadiraju, and M. K. Medam, "Satellite-based AIS receiver for dense maritime zones," in Proc. 9th Int. Conf. on Communication Systems and Networks (COMSNETS), 2017, pp. 15–22.

[6] M. Picard, M. R. Oularbi, G. Flandin, and S. Houcke, "An adaptive multi-user multi-antenna receiver for satellite-based AIS detection," in Proc. 6th Advanced Satellite Multimedia Systems Conf. (ASMS) and 12th Signal Processing for Space Communications Workshop (SPSC), 2012, pp. 273–280.

[7] S. Uehashi, Y. Nouda, M. Hangai, and T. Ohgane, "Successive Interference Cancellation for Asynchronous Signal Collision in Space-Based AIS," in Proc. IEEE Aerospace Conf. (AERO), Big Sky, MT, USA, Mar. 2022, pp. 1–7, doi: 10.1109/AERO53065.2022.9843214.

[8] G. Wang, L. Li, X. Chen, and Z. Zhang, "Deep learning-based AIS signal collision detection in a satellite reception environment," Applied Sciences (MDPI), vol. 16, no. 2, Art. 643, 2026. [Online]. Available: https://www.mdpi.com/2076-3417/16/2/6

[9] Honeywell International Inc., "Satellite-AIS Receiver AIS-MS03," Honeywell Aerospace, Phoenix, AZ, USA, Datasheet N61-1899-000-000, Mar. 2018. Accessed: Sep. 4, 2026. [Online]. Available: https://research.csiro.au/cceo/wp-content/uploads/sites/252/2020/09/NovaSAR-1_AIS_Receiver.pdf

[10] M. J. Gallardo and U. Sorger, "Coherent receiver for AIS satellite detection," in Proc. 4th Int. Symp. on Communications, Control and Signal Processing (ISCCSP), 2010, pp. 1–4.

[11] X. Meng, C. Liu, J. Teng, et al., "Frequency offset estimation of satellite-based AIS signals based on interpolated FFT," Wireless Personal Communications, vol. 99, no. 1, pp. 35–45, 2018.

[12] M. Morelli and U. Mengali, "Joint frequency and timing recovery for MSK-type modulation," IEEE Trans. Commun., vol. 47, no. 6, pp. 938–946, 1999.

[13] D. C. Rife and G. A. Vincent, "Use of the discrete Fourier transform in the measurement of frequencies and levels of tones," Bell Syst. Tech. J., vol. 49, no. 2, pp. 197–228, Feb. 1970.

[14] D. J. Nelson, J. R. Hopkins, and A. L. Bartos, "Coherent demodulation of AIS-GMSK signals in co-channel interference," in Proc. 45th Asilomar Conf. Signals, Systems and Computers, 2011, pp. 295–299.

[15] Recommendation ITU-R M.1371 (Antenna characteristics), technical annex on ship dipole antenna radiation pattern.

[16] Anthropic, "The Claude 3 model family: Opus, Sonnet, Haiku," Anthropic, San Francisco, CA, USA, Tech. Rep., Mar. 2024. Accessed: Sep. 1, 2026. [Online]. Available: https://www-cdn.anthropic.com/de8ba9b01c9ab7cbabf5c33b80b7bbc618857627/Model_Card_Claude_3.pdf

[17] Kpler MarineTraffic, "Live Ships Map — Global AIS vessel density visualization," AIS Marine Traffic. Accessed: Sep. 1, 2026. [Online]. Available: https://www.marinetraffic.com/en/ais/home

[18] A. Macikunas, J. Safar, and N. Ward, "VDES Channel Model - Satellite Channel Characteristics," IALA ENAV20 Meeting, Doc. D1.17A, France, Mar. 2017. Accessed: Sep. 1, 2026. [Online]. Available: https://www.iala.int/content/uploads/2018/05/D1.17A-Presentation-on-VDES-satellite-test-results-to-ENAV20-WG3-V3.pdf