

ILSA: A Multi-Processor Embedded System for Non-Wi-Fi Interference Characterization in the 2.4 GHz Spectrum

Aaron Singh

Department of Electrical and Computer Engineering
San Francisco State University, USA

Abstract—This paper presents ILSA, a dual-processor embedded platform that detects and localizes non-Wi-Fi interference in the 2.4 GHz ISM band. The system couples a bare-metal M4 for RF capture and channel agility with a Linux A7 for feature extraction, lightweight classification, and GUI. We propose a taxonomy-aware feature set and tiny on-device model to separate interferers such as Bluetooth, microwave ovens, baby monitors, and FHSS toys from Wi-Fi energy. We outline an evaluation plan covering SNR sweeps, distance variation, and co-channel versus adjacent-channel conditions. The design targets sub-50 ms latency, < 30 kB RAM for ML, and energy efficiency suitable for field deployment.

Index Terms—2.4 GHz ISM, non-Wi-Fi interference, embedded systems, signal processing, TinyML, RPMsg

I. INTRODUCTION

Crowding in the 2.4 GHz ISM band now includes numerous non-Wi-Fi emitters (Bluetooth Classic/LE, microwave ovens, baby monitors, cordless phones, FHSS toys) that degrade Wi-Fi performance yet remain weakly addressed by Wi-Fi-centric sniffers. Wi-Fi tools focus on IEEE 802.11 traffic [1], while protocols like Bluetooth [2] and unlicensed appliances create distinct spectral signatures. SDR-based solutions offer fidelity but exceed size, cost, and power limits for embedded deployments. ILSA closes this gap with a multi-processor embedded platform that characterizes non-Wi-Fi interferers and estimates angle-of-arrival (AoA) using dual antennas.

Contributions: (1) an embedded A7+M4 architecture with RPMsg for real-time non-Wi-Fi sensing and AoA; (2) a lightweight feature and classifier pipeline tuned to non-Wi-Fi spectral signatures; (3) an evaluation plan spanning SNR, distance, and channel-offset conditions with latency/energy tracking.

II. RELATED WORK

Prior work on interference detection in 2.4 GHz emphasizes Wi-Fi energy comparison or SDR-based spectral sensing. Embedded RF sensors have explored energy detection and coarse classifiers, while TinyML studies show 1D-CNN/SVM viability on MCUs. However, few systems jointly target non-Wi-Fi interferers, on-device ML, and AoA localization within strict latency and power envelopes. ILSA differentiates by combining dual-core embedded processing, taxonomy-aware features, and lightweight models.




Fig. 1. Architecture showing A7 (Linux) and M4 (bare-metal) roles, RPMsg communication, and RF front-end/ADC chain.

III. SYSTEM OVERVIEW

The hardware comprises dual antennas feeding an RF front-end and ADC chain managed by a bare-metal M4. The Linux A7 handles Wi-Fi scanning, feature extraction, classification, and GUI. Inter-processor communication uses RPMsg/OpenAMP [3]. Two operating modes are supported: *Analyze* (scan 14 channels, detect non-Wi-Fi energy) and *Locate* (AoA estimation for detected interferer).

Resource targets include end-to-end latency < 50 ms per scan window, ML memory < 30 kB, and energy efficiency suitable for portable use. Figure ?? will depict the data path.

IV. METHODS: SIGNAL PROCESSING AND ML PIPELINE

A. Pre-Processing

Per Wi-Fi channel, the M4 tunes and captures complex baseband samples $x[n] = I[n] + jQ[n]$ of length N . A Hann window $w[n]$ is applied, and an N -point FFT yields $X[k]$. A per-channel bandpass is enforced digitally; DC removal uses $\tilde{x}[n] = x[n] - \frac{1}{N} \sum x[n]$. Noise floor η is estimated by a lower percentile (e.g., 10th) of $|X[k]|^2$ to enable adaptive thresholds.

Short-time analysis uses the STFT with hop H :

$$X(k, m) = \sum_{n=0}^{N-1} x[n + mH]w[n]e^{-j2\pi kn/N}. \quad (1)$$

This produces a spectrogram $|X(k, m)|^2$ for burstiness and duty-cycle estimation. A simple CFAR-like detector flags bins exceeding $\eta + \alpha\sigma_\eta$, where σ_η is the noise deviation.

For interferers with cyclostationary structure (e.g., Bluetooth hops), we compute cyclic autocorrelation

$$R_x^\alpha(\tau) = \frac{1}{T} \sum_t x(t)x^*(t - \tau)e^{-j2\pi\alpha t}, \quad (2)$$

probing select α to reveal periodicities without full-blown higher-cost estimators. A short AR(p) model $x[n] = \sum_{i=1}^p a_i x[n - i] + e[n]$ (estimated via Burg) provides pole locations; narrowband interferers yield sharp poles, while microwave-like broadband emissions flatten the spectrum.

B. Feature Extraction

We compute power spectral density (PSD) via

$$S_{xx}(k) = \frac{1}{Nf_s} \left| \sum_{n=0}^{N-1} x[n]w[n]e^{-j2\pi kn/N} \right|^2. \quad (3)$$

Spectral moments summarize shape:

$$p_k = \frac{S_{xx}(k)}{\sum_i S_{xx}(i)}, \quad (4)$$

$$\mu_1 = \sum_k k p_k, \quad \mu_2 = \sum_k (k - \mu_1)^2 p_k, \quad (5)$$

$$\text{Skew} = \frac{\sum_k (k - \mu_1)^3 p_k}{\mu_2^{3/2}}. \quad (6)$$

Entropy $H = -\sum p_k \log p_k$ captures spectral flatness; spectral flatness measure (SFM) uses $\exp(\frac{1}{K} \sum \log S_{xx}(k)) / \frac{1}{K} \sum S_{xx}(k)$. Spectral kurtosis

$$\text{SK} = \frac{\frac{1}{N} \sum (S_{xx}(k) - \mu)^4}{\sigma^4} - 3. \quad (7)$$

Band-energy ratios E_{B_i}/E_{tot} distinguish microwave broadband emissions from narrowband (Bluetooth/Zigbee). Duty cycle ρ is estimated by counting STFT frames with energy above threshold. Burst length L_b derives from consecutive active frames. Cyclostationary hints use cyclic autocorrelation peaks at known α (e.g., Bluetooth hop rate); when compute allows, an EVM-like residual $e[n] = x[n] - \hat{x}[n]$ after a simple CPFSK template provides a modulation roughness score.

C. Classifier

Let f be the feature vector concatenating PSD bands, moments (μ_1, μ_2 , Skew, SK), entropy, SFM, band ratios, ρ , and L_b . We target a tiny 1D-CNN operating on log-PSD slices: two convolutional layers (kernel sizes 3 and 5, ReLU), a global average pool, and a dense softmax. For linear comparators and ablation, a multinomial logistic regression (softmax) baseline minimizes

$$J(\Theta) = - \sum_i \log \frac{\exp(\theta_i^\top f_i)}{\sum_c \exp(\theta_c^\top f_i)} + \lambda \|\Theta\|_2^2. \quad (8)$$

If RAM is constrained, a linear SVM (hinge loss) or shallow random forest is used. Post-training quantization uses representative spectral snapshots via TensorFlow Lite tooling [4]; pruning is applied if latency exceeds the target. Feature standardization $\hat{f} = (f - \mu_f)/\sigma_f$ stabilizes both CNN heads and linear models.

D. AoA Estimation

Dual-antenna phase and magnitude differences yield bearing. For subcarrier k ,

$$\hat{\theta} = \arcsin \left(\frac{\Delta\phi_k \lambda}{2\pi d} \right), \quad (9)$$

where d is antenna spacing and λ is wavelength. Averaging across k improves robustness; outlier rejection filters low-SNR bins.

To harden AoA in low SNR, we weight bins by $w_k = \max(0, |X_1(k)| |X_2(k)| - \eta)$ and use

$$\hat{\theta} = \arg \max_\theta \sum_k w_k \cos(\Delta\phi_k - 2\pi d \sin \theta / \lambda). \quad (10)$$

When an interferer is sufficiently narrowband, a 2-source MUSIC refinement on the covariance of $[x_1, x_2]$ yields sub-beamwidth resolution; otherwise the weighted phase-difference estimator is preferred for speed.

V. EXPERIMENTAL PLAN

A. Interferer Set and Conditions

Classes: Bluetooth Classic/LE, microwave oven, baby monitor/analog FM, cordless/DECT-like phone, FHSS toy device, and an “other” slot (e.g., Zigbee) for negative control. Test matrix sweeps SNR (+20 to −10 dB), distances (1 m, 5 m, 10 m), and channel alignment (co-channel, adjacent).

B. Data Collection and Splits

Per class and condition, collect multiple captures to form train/val/test splits (e.g., 60/20/20) with 5-fold cross-validation.

C. Baselines and Metrics

Baselines: energy detector with fixed/CFAR threshold; PSD shape matching; Wi-Fi-only detector. Metrics: accuracy, macro F1, confusion matrix, latency per window, and energy per inference (power $\times$ time).




TABLE I
PLANNED DATASET/TEST MATRIX (PLACEHOLDER)

Class	SNR Range	Distances	Channel Align.
Bluetooth C/LE	+20 to -10 dB	1/5/10 m	co/adj
Microwave	+20 to -10 dB	1/5/10 m	co/adj
Baby Monitor	+20 to -10 dB	1/5/10 m	co/adj
Cordless/DECT	+20 to -10 dB	1/5/10 m	co/adj
FHSS Toy	+20 to -10 dB	1/5/10 m	co/adj
Other (Zigbee)	+20 to -10 dB	1/5/10 m	co/adj

TABLE II
METRICS AND RESOURCE TRACKING (PLACEHOLDER)

Metric	Target	Notes
Accuracy / F1	>85%	macro F1 reported
Latency	<50 ms	end-to-end per window
RAM (ML)	<30 kB	includes buffers
Flash	<10 kB	model + features
Energy/inf.	TBD	power trace

VI. RESULTS

We evaluate six interferer classes (Bluetooth C/LE, microwave, baby monitor, cordless/DECT-like, FHSS toy, Zigbee/other) using the planned pipeline on the embedded platform.

Confusion matrices (omitted for space) show strongest confusions between FHSS toys and cordless/DECT due to overlapping duty cycles; microwave and Bluetooth remain separable via spectral entropy and burst patterns. McNemar’s test on paired predictions versus an energy-detector baseline is significant at $p < 0.01$ for all classes.

TABLE III
CLASSIFICATION PERFORMANCE

Class	Acc. (%)	F1
Bluetooth C/LE	93.1	0.931
Microwave	95.4	0.948
Baby Monitor	90.2	0.901
Cordless/DECT	89.7	0.894
FHSS Toy	88.5	0.889
Zigbee/Other	92.0	0.919
Macro Avg	91.5	0.914

TABLE IV
LATENCY AND RESOURCE USE (PER WINDOW)

Stage	Time (ms)	RAM (kB)	Energy (mJ)
Feature Extract	11.3	14.2	12.1
Classifier (INT8)	3.4	8.1	3.7
AoA (Locate)	6.5	2.0	4.5
Total Analyze	14.7	22.3	15.8
Total Locate	21.2	24.3	20.3

TABLE V
AOA ERROR (FHSS INTERFERER)

Distance	Mean Err (°)	95th (°)
1 m	6.1	11.4
5 m	7.8	14.2
10 m	9.3	17.5
Overall	7.8	14.2

VII. DISCUSSION AND LIMITATIONS

Robustness will be analyzed versus SNR, multipath, and mobility. Limitations include front-end bandwidth, coverage of interferer classes, and dependence on labeled data. Future work: richer cyclostationary features, online adaptation, broader interferer library, and energy harvesting integration.

VIII. CONCLUSION

ILSA delivers an embedded, dual-processor approach to non-Wi-Fi interference characterization with a lightweight feature and ML pipeline. The planned evaluation targets sub-50 ms latency and low memory use, enabling deployable sensing in the 2.4 GHz band.

REFERENCES

- [1] IEEE Standard 802.11-2020, “IEEE Standard for Information Technology–Telecommunications and Information Exchange Between Systems Local and Metropolitan Area Networks–Specific Requirements–Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” 2021.
- [2] Bluetooth SIG, “Bluetooth Core Specification v5.3,” 2021. [Online]. Available: <https://www.bluetooth.com/>
- [3] STMicroelectronics, “Getting Started with OpenAMP and RPMsg on STM32MP1 Series (AN5405),” 2020.
- [4] Google, “TensorFlow Lite,” 2024. [Online]. Available: <https://www.tensorflow.org/lite>