Spectrogram-Based Interference Source Attribution at OVRO: An Over-the-Air Validation Using Pseudonymetry

Meles G. Weldegebriel*, Zihan Li†, Greg Hellbourg‡, Ning Zhang†, Neal Patwari§

*Electrical and Systems Engineering, Washington University in St. Louis, USA
†Computer Science and Engineering, Washington University in St. Louis, USA
‡California Institute of Technology, USA
§University of Utah, USA

***Abstract*—The operation of highly sensitive passive receivers such as radio telescopes can be disrupted even when the radio frequency interference (RFI) is below the noise floor. Protecting these receivers and attributing the interference to a specific transmitter is particularly challenging. We present an over-the-air field validation of spectrogram-based interference source attribution using radio telescope backend spectrogram data. The validation is performed at the Owens Valley Radio Observatory (OVRO) with the DSA-110 system. The secondary user embeds a pseudonym watermark on a single OFDM subchannel. The passive receiver recovers the pseudonym directly from the scientific spectrogram data without additional RF hardware. We design a DSA-110 backend-compatible pseudonym receiver pipeline enabling blind synchronization, timing compensation, and low-SNR pseudonym decoding. Measured bit error performance with -15 to -5 dB per-bin SNR demonstrates the feasibility of spectrogram-based attribution with performance comparable to an IQ-based baseline.**

***Index Terms*—Radio Frequency Interference, Passive Receiver, Radio Astronomy, Pseudonymetry, Interference Mitigation.**

I. Introduction

Spectrum coexistence between active wireless systems and passive systems is increasingly challenged by radio frequency interference (RFI). Radio telescopes are particularly vulnerable, as they are designed to detect extremely weak signals from distant objects. Coexistence requires reliable detection and transmitter attribution at very low signal-to-noise ratio (SNR).

Traditional RFI management strategies include wide geographic quiet zones and propagation-model based planning [1]. When interference is detected, passive systems use reactive mitigation techniques such as filtering, excision, and statistical flagging [2]–[4]. However, these methods have limitations. Quiet zones are often overly conservative, leading to spectrum inefficiency, and generally not effective for spaceborne interference. In addition, post-processing cannot always recover lost data especially when interference is weak, intermittent, or aggregated.

Recent research has proposed more flexible coexistence strategies, including cooperative mitigation [5], dynamic protection areas (DPAs) [6], [7], and machine learning for RFI detection [8]. These techniques introduce adaptive spectrum management. Yet a critical gap remains: even when RFI is detected, it is often impossible to identify the specific interfering device. Without device-level accountability, such coexistence frameworks resort to conservative assumptions.

In this work, our primary focus is detecting an interference event and attributing it to a particular interferer using radio astronomy receiver measurements. Our approach is based on pseudonymetry, in which identification of an interfering transmitter triggers automated reassignment or shutdown [9], [10]. In this cooperative method, a secondary transmitter watermarks its signal transmission with a lightweight pseudonym. The watermark is designed for reliable detection and transmitter identification under low-SNR conditions.

A key challenge in deploying pseudonymetry at real observatories is receiver integration. Radio telescope backends are designed for scientific data products and typically output spectrograms or correlated visibility data [11], [12]. Additionally, practical deployment must address the inherent mismatches between interference source and backend receiver parameters including bandwidth, frequency resolution, and timing resolution (Table I). This motivates the following question: *Can pseudonym watermarks be extracted using spectrogram outputs from standard radio telescope receivers, despite these parameter mismatches and lack of phase information?*

We investigate this question through an over-the-air field trial at the Owens Valley Radio Observatory (OVRO), using the Deep Synoptic Array (DSA-110) digital backend [11], [12]. An SDR-based OFDM transmitter operating outside the telescope's observing band serves as controlled interference source and embeds a pseudonym watermark on a single OFDM subcarrier. The DSA-110 backend produces beamformed spectrograms and cross-correlated antenna measurements. We develop and evaluate a pseudonym receiver pipeline that performs time-correlation-based packet alignment, resampling to correct timing mismatches, and energy-domain pattern matching for watermark decoding.

The main contributions of this work can be summarized as:

- The first validation of spectrogram-based pseudonym decoding on an operational radio astronomy backend, demonstrating transmitter attribution without access to baseband IQ samples.
- A backend-compatible receiver design with parameter

mismatch compensation, enabling blind synchronization and reliable low-SNR decoding under realistic telescope resolution constraints.
- An experimental evaluation of attribution performance from -15 dB to -5 dB per-bin SNR, showing comparable performance to an IQ-based baseline.

II. Background and Related Work

A. RFI Detection and Protection

For passive services like radio astronomy, protection criteria are particularly challenging. The ITU-R RA.769 recommendation defines extremely low harmful interference thresholds, emphasizing that even extremely weak emissions that are far below the reception thresholds of conventional communication systems can corrupt radio telescope observations [13]. Radio telescope receivers use a combination of site protection, front-end filtering, and post-processing techniques to mitigate interference impact. Adaptive thresholding, excision, and flagging are widely used [2]–[4]. The LOFAR telescope, for instance, has demonstrated a practical mitigation pipeline with strong real-world results [4].

However, most methods are reactive, i.e., they mitigate interference after it has already occurred. Moreover, they lack device-level accountability capabilities — when RFI is detected, there is typically no mechanism to identify the specific interfering device. As RFI becomes more frequent, the above reactive approaches become increasingly expensive, time-consuming, and less effective. This motivates the need for cooperative and accountable approaches that integrate monitoring and real-time interference management.

Coexistence between active and passive services has been studied in various contexts. Some work has examined interference mitigation for satellite systems sharing spectrum with passive receivers [5]. More recently, dynamic protection area (DPA) methods that aim to manage sharing more flexibly have been proposed [6]. These approaches can improve spectrum utilization [7], but without reliable interference detection, identification, and real-time control, these dynamic protection frameworks tend to default to the traditional conservative approaches.

Machine learning techniques, such as deep learning, have also been proposed for RFI detection in radio astronomy data [8]. These methods enable identification of complex interference patterns and support automation. However, detection alone does not provide device-level attribution. A passive receiver may know that RFI exists, but if the specific interfering device is not known, real-time interference management becomes insufficient. Our work complements ML-based detection where machine learning can identify interference events, and our watermarking-based approach enables transmitter identity.

B. Cooperative Attribution and Pseudonymetry

Pseudonymetry has been proposed as a cooperative coexistence mechanism that enables transmitter attribution through embedded signal watermarks [9]. A transmitter inserts a lightweight identifier into its waveform in a way that allows energy-based reception, enabling recovery by passive receivers even when conventional phase-based demodulation is infeasible. A key advantage is its ability to operate at very low SNR levels. This makes pseudonymetry well suited to radio astronomy environments, where harmful interference occurs far below communication receiver sensitivity thresholds [13].

However, previous studies evaluated pseudonymetry using receivers with access to baseband IQ samples. In practical radio astronomy systems, digital backends typically provide energy data in the time, frequency, or space domains, instead of raw IQ samples, since multi-antenna wideband sampling would generate IQ data rates impractical for routine operations. This work therefore evaluates pseudonym watermark recovery using spectrogram data to validate its performance in real-world passive receiver systems.

III. System Architecture and Watermark Model

A. System Architecture

We consider a cooperative spectrum-sharing scenario with three components: a secondary transmitter, an operational radio astronomy backend, and a backend-compatible pseudonym receiver. The goal is to detect when interference happens and attribute it to a specific cooperative transmitter, using only the standard outputs provided by the telescope backend.

Modern radio astronomy backends typically export channelized spectra or visibilities rather than continuous baseband IQ streams. Continuous baseband export from multiple antennas at multi-GHz bandwidth would result in data rates on the order of terabits per second, which are incompatible with routine storage and network transport in operational observatories. Introducing parallel SDR receivers or modifying telescope front-ends would increase deployment cost and complexity. Therefore, evaluating pseudonym recovery using standard beamformed spectrogram outputs, without adding dedicated receiver hardware, is essential for practical deployment.

The architecture assumes that interference may arrive through non-ideal spatial paths, including antenna sidelobes or indirect propagation mechanisms. Therefore, the pseudonym receiver must operate under weak-signal conditions without assuming main-lobe alignment or favorable geometry. This constraint motivates energy-domain processing and low-SNR synchronization capability.

The system consists of three functional blocks:
- **Secondary transmitter:** an active device embedding a pseudonym watermark in its waveform.
- **Passive receiver:** a radio astronomy backend exporting beamformed spectrogram data.
- **Pseudonym receiver:** a backend-compatible post-processing pipeline that extracts the pseudonym.

The passive receiver has no prior knowledge of transmission start time and must perform blind synchronization. It only observes energy within a wide channelized bandwidth. The RFI appears as narrowband energy inside the spectrogram.

We design the system to match backend limitations. In practice, radio astronomy backends generate spectrograms for

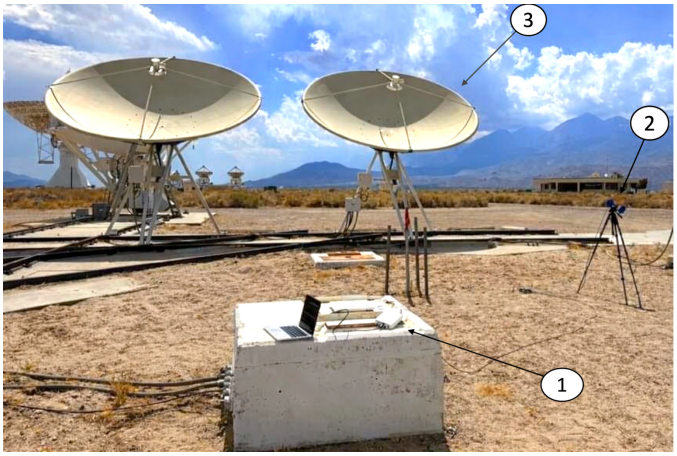

Fig. 1: Field trial setup at OVRO. (1) USRP B210 with laptop control, (2) Bicolog 20300 transmit antenna, (3) DSA-110 receiving antenna.

scientific use. If pseudonym receiver pipeline can use this same data, the deployment cost is reduced. Also, spectrogram-based decoding means the approach can be applied retroactively to stored data, enabling offline attribution analysis of interference events. This supports accountability and coordination within dynamic protection frameworks.

B. Single-Subcarrier Coded Pulse Amplitude Modulation

The pseudonym watermark is embedded using coded pulse amplitude modulation (CPAM) [10], in which each pseudonym bit is represented by a spreading sequence of L chips with predefined amplitude pattern. This spreading structure distributes each bit over multiple chip intervals and enables reliable detection at low SNR even with a changing channel.

Let $f(t)$ denote the sinusoidal signal on the pseudonym subcarrier, and let $Q_p(t)$ denote the coded watermark corresponding to pseudonym bit $p \in \{0, 1\}$. The watermark signal is defined as

$$Q_p(t) = \sum_{l=0}^{L-1} (1 - \alpha A_p[l])\ \phi(t - lT_c), \tag{1}$$

where $A_p[l]$ is the chip amplitude sequence, α is the modulation index, T_c is the chip duration, and $\phi(t)$ is the chip pulse shape. The transmitted watermarked signal in the pseudonym subcarrier is

$$s_p(t) = Q_p(t)\, f(t).$$

In this work, the chip amplitude sequence is generated using a maximum-length pseudo-noise (PN) sequence of length $L = 15$. Pseudonym bit '1' is represented by the chip sequence, referred to here as patterns, $A_1[l]$, where $l \in \{0, \ldots, L-1\}$ indexes the chip position within one bit. Bit '0' is represented by its complement $A_0[l] = 1 - A_1[l]$. These sequences determine the chip amplitude pattern in coded modulation.

C. Energy-domain Watermark Representation

Chip values modulate subcarrier power: chip '0' corresponds to reduced power and chip '1' to increased power relative to nominal. Therefore, the watermark appears as a structured power variation pattern in the extracted spectrogram time series, as shown in Fig. 2.

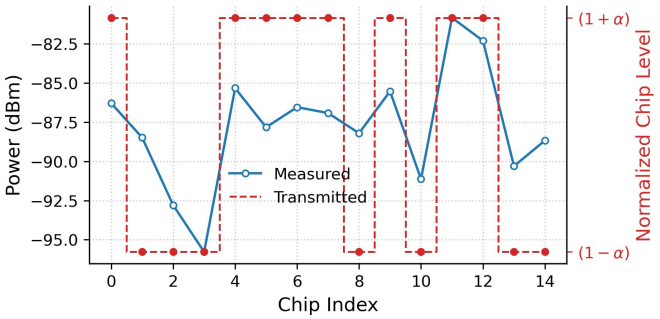

Fig. 2: Transmitted pseudonym power pattern (-●-), and experimental received power (–○–) from OVRO/DSA-110 backend.

This watermark design fits our design requirements because it does not depend on coherent demodulation – measured power changes in time reveal the watermark. Also, using a single subcarrier minimizes overhead and complexity: the pseudonym receiver needs to extract only one frequency subchannel from the telescope spectrogram.

IV. Spectrogram-Based Receiver

A. Spectrogram Processing

A model of the radio telescope and pseudonym signal chain is shown in Fig. 3. The received RF signal is first processed by the RF front-end, which includes the feed, low-noise amplifier (LNA), filter, and fiber transmitter. The signal is sent to the DSA-110 back-end receiver via fiber. It is amplified and filtered before digital signal processing.

The telescope digital receiver performs analog-to-digital conversion (ADC), followed by channelization, equalization, and requantization. The resulting channelized samples are then packetized and sent over the network to a GPU server responsible for beamforming. Beamformed spectrograms, representing power distribution in the time and frequency domains for one given direction within the telescope primary beam, are then captured and processed to detect RFI and search for astrophysical transient signals. The pseudonym receiver operates directly on these spectrograms. Since the processing is performed in the energy domain, the receiver cannot perform carrier recovery, phase tracking, or symbol-level demodulation. Instead, it analyzes structured temporal and spectral energy variations introduced by the pseudonym watermark.

The pseudonym receiver pipeline consists of three main stages. First, the synchronization stage applies spectral correlation to detect the presence of a watermark and estimate packet timing. Second, the resolution alignment stage applies resampling to mitigate the resolution difference between the transmitter symbol duration and the spectrogram integration window. Finally, the pseudonym extraction stage recovers the embedded pseudonym bits by comparing measured energy patterns in the selected frequency bins with the transmitted bit patterns.

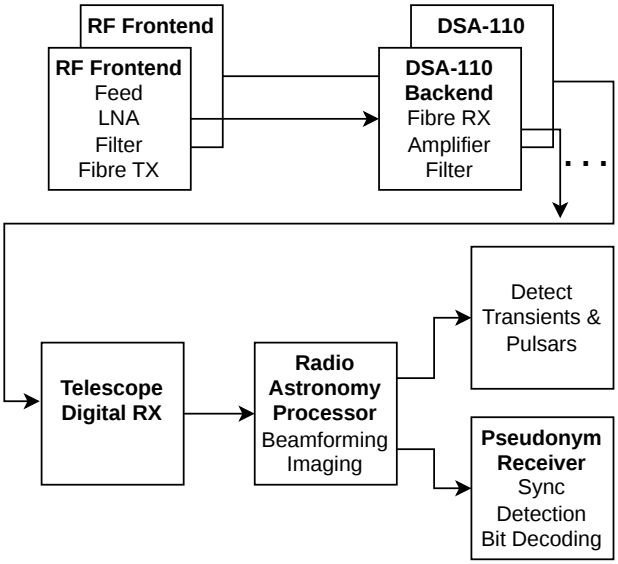

Fig. 3: Signal processing chain and pseudonym receiver pipeline integrated with the OVRO DSA-110 digital receiver.

B. Synchronization and Detection

Pseudonym packet detection and synchronization operate via cross-correlation between the received power stream and a known transmitted reference CPAM watermark. High correlation indicates detection of the watermark, and the correlation peak gives an estimated start time.

One critical challenge in this field trial is the inherent difference in time and frequency resolutions between the secondary transmitter and the radio telescope backend. Because of this, the transmitted watermark chip pattern does not perfectly align with the received spectrogram samples.

1) Passive receiver: The OVRO/DSA-110 backend scans a wide bandwidth (about 2 GHz) and produces *spectrogram* (power) data, not IQ samples. During our experiment, we use 11 μs time resolution and 90 kHz frequency resolution. Therefore, we can only decode the pseudonym by using the power time series from the corresponding spectrogram frequency bin.

2) Timing mismatch: The transmitter watermark timing is defined by OFDM parameters, while receiver timing is determined by spectrometer integration time. This creates a small but important mismatch. In our experiments, these mismatches are $T_{\mathrm{TX}} = (93.75\ \mathrm{kHz})^{-1} \approx 10.67\ \mu$s and $T_{\mathrm{RX}} = (90\ \mathrm{kHz})^{-1} \approx 11.11\ \mu$s. Although small, this timing mismatch accumulates over multiple chips and, without correction, chip boundaries drift relative to spectrogram samples and correlation alignment is degraded.

To compensate, we resample the extracted spectrogram power stream by a ratio of T_{RX}/T_{TX} (via interpolation) and then oversample by a factor of 10 to stabilize chip-boundary alignment. After alignment, chip-level averaging reduces noise variance and stabilizes decoding under low-SNR conditions.

3) Bit length alignment: One pseudonym bit corresponds to $N_{\mathrm{TX}} = T_b/T_{\mathrm{TX}} = 960\ \mu\mathrm{s}/10.67\ \mu\mathrm{s} \approx 90$ transmitter samples. Due to backend resolution differences, the receiver observes approximately $N_{\mathrm{RX}} \approx 86.4$. After resampling and oversampling, each pseudonym bit spans roughly 900 samples, yielding 60 samples per chip. Averaging within each chip interval reduces noise variance and improves decoding robustness.

4) Resolution constraints: Pseudonym detectability in a spectrogram-based data outputs is governed not only by received power, but also by backend frequency and time resolution.

a) Frequency resolution: Let Δf_{rx} denote the backend spectral resolution (i.e., the equivalent noise bandwidth of one analysis channel), and let B_{sc} denote the effective bandwidth of the pseudonym subcarrier centered at f_0. Let P_{sc} denote the average received pseudonym subcarrier power, and let N_0 denote the one-sided noise power spectral density (W/Hz).

Only the portion of the pseudonym spectrum overlapping a backend channel contributes to detection. Define the overlap efficiency $\eta_f \in [0, 1]$ as the fraction of pseudonym energy captured by a given backend bin. The signal-to-noise ratio in a single spectrogram bin is therefore approximately

$$\mathrm{SNR}_{\mathrm{bin}} \approx \frac{P_{\mathrm{sc}}\,\eta_f}{N_0\,\Delta f_{\mathrm{rx}}}, \tag{2}$$

where $\mathrm{SNR}_{\mathrm{bin}}$ denotes the per-bin signal-to-noise ratio.

If $\Delta f_{\mathrm{rx}} \gtrsim B_{\mathrm{sc}}$, most pseudonym energy is captured in a single bin ($\eta_f \approx 1$). If $\Delta f_{\mathrm{rx}} \lesssim B_{\mathrm{sc}}$, the pseudonym energy is distributed across multiple adjacent bins and must be combined to avoid detection loss. In this regime, incoherent aggregation across the occupied bins recovers the full signal energy, yielding an aggregate SNR,

$$\mathrm{SNR}_{\mathrm{agg}} \approx \frac{P_{\mathrm{sc}}}{N_0\,\Delta f_{\mathrm{rx}}}, \tag{3}$$

thus detectability depends on the total integrated bandwidth rather than how that bandwidth is partitioned into individual channels.

For OVRO, $\Delta f_{\mathrm{rx}} = 90$ kHz and $B_{\mathrm{sc}} = \Delta f = 93.75$ kHz, placing the system close to the transition between single-bin capture and multi-bin energy spreading. Most of the pseudonym energy therefore falls within one backend bin, preserving effective SNR.

b) Time resolution.: Let R_b be the pseudonym bit rate and bit duration $T_b = \frac{1}{R_b} = LT_c$. The maximum number of samples that can be averaged within one bit interval is

$$N_{\max} = \left\lfloor \frac{T_b}{T_{\mathrm{RX}}} \right\rfloor. \tag{4}$$

Since incoherent (power) averaging reduces noise variance proportionally to $1/N$, averaging $N_{\max}$ independent samples increases the effective bit-level SNR approximately by a factor of $N_{\max}$. Substituting $\mathrm{SNR}_{\mathrm{bin}}$ yields

$$\mathrm{SNR}_{\mathrm{bit}} \approx \frac{P_{\mathrm{sc}}\,\eta_f}{N_0\,\Delta f_{\mathrm{rx}}} \cdot N_{\max}, \tag{5}$$

where $\mathrm{SNR}_{\mathrm{bit}}$ denotes the SNR available for detecting one pseudonym bit.

Algorithm 1 Spectrogram-based pseudonym decoding

Require: Spectrogram $S(t, f)$, watermark bin f_0, patterns $A_0[l], A_1[l]$
Ensure: Decoded bits $\hat{b}$
1: Extract $x[t] \leftarrow S(t, f_0)$
2: Start $\tau \leftarrow \arg\max \mathrm{corr}(x, \mathrm{template})$
3: Align $x[t] \leftarrow x[t + \tau]$
4: Resample $x \leftarrow \mathrm{Resample}(x, Resample(x, T_{\mathrm{RX}}/T_{\mathrm{TX}}))$
5: Average to chip energies $z[k] \leftarrow \mathrm{AvgChip}(x)$
6: Decode by matching z with A_0, A_1

In the OVRO experiment, $T_b \approx 960\,\mu$s and $T_{\mathrm{RX}} = 11\,\mu$s, yielding approximately 86 samples per bit before resampling. This integration gain explains the low bit-error probabilities observed for SNR above approximately -8 dB.

c) Rate–integration trade-off: Let γ denote the minimum SNR required for reliable bit detection. Since the maximum number of independent backend samples per bit is

$$N_{\max} \approx \frac{T_b}{T_{\mathrm{RX}}} = \frac{1}{R_b\, T_{\mathrm{RX}}},$$

Using bit-level SNR in (5),the reliability condition $\mathrm{SNR}_{\mathrm{bit}} \geq \gamma$ therefore implies

$$R_{b,\max} \approx \frac{P_{\mathrm{sc}}\, \eta_f}{\gamma\, N_0\, \Delta f_{\mathrm{rx}}\, T_{\mathrm{RX}}}. \tag{6}$$

Thus, passive spectrogram-based receivers impose a rate–integration trade-off rather than a strict power threshold: weaker pseudonym signals remain attributable if transmitted more slowly, allowing sufficient time averaging within each bit.

C. Pseudonym Decoding

Decoding compares the averaged chip-energy sequence to the known patterns $A_0[l]$ and $A_1[l]$. The main computation is correlation and averaging. Since this uses only one frequency bin, the cost is low compared to full-band ML approaches [8]. This suggests the receiver can be deployed as a lightweight post-processing step for accountability and coexistence reporting.

V. Experimental Setup and Results

A. Experimental Setup

Transmitter. A USRP B210 was used as the cooperative secondary transmitter. It was connected to a Bicolog 20300 broadband antenna (20 MHz–3 GHz). The transmitter bandwidth is $f_s = 6$ MHz with $N = 64$ subcarriers, resulting in subcarrier spacing $\Delta f = 93.75$ kHz. Transmitter gain varied from 10 dB to 32 dB. As discussed in Section III, we use one subcarrier for the watermark communication and we refer to it as the *pseudonym subcarrier*. The chip duration is 64 µs, which corresponds to 6 symbol periods. Therefore, each pseudonym bit occupies 960 µs, and a pseudonym packet containing 28 bits spans 26.90 ms.

TABLE I: TX and RX Parameters

Parameter	Symbol	Value (Experiment)
Communication Signal		
OFDM bandwidth	f_s	6 MHz
Number of subcarriers	N	64
Subcarrier spacing	Δf	93.75 kHz
Pseudonym subcarrier bandwidth	B_{sc}	93.75 kHz
Pseudonym Transmission		
Watermark subcarriers	N_w	1
Bit rate	R_b	1.05 kbit/s
Bit duration	T_b	960 μs
Chips per bit	L	15
Radio Astronomy Receiver		
Backend bandwidth	B	76.77 MHz
Frequency resolution	Δf_{rx}	90 kHz
Time resolution	T_{RX}	11 μs

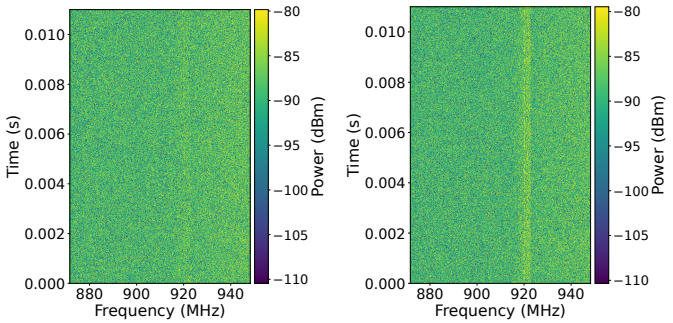

(a) -5 dB (b) $+5$ dB

Fig. 4: Spectrogram examples at two SNR levels.

Receiver. The passive receiver is the OVRO/DSA-110 backend spectrometer. Although each antenna is digitized at high rate, only a selected sub-band is channelized and exported for astronomical processing. The operation band is sampled at a 90 kHz resolution. The spectrometer time resolution is 11 μs.

B. Decoding Performance

We decode pseudonym bits from spectrogram measurements and compute probability of pseudonym bit error P_e vs. SNR. SNR is computed per spectrogram frequency bin at the backend output. This empirical definition corresponds to the per-bin SNR derived in (2), where signal power is estimated from measured spectrogram values rather than analytical power spectral density (PSD) quantities. Noise power is measured in the watermark bin when the transmitter is OFF, and signal-plus-noise power is measured at the same bin when the transmitter is ON. Signal power is estimated as the linear difference, and $\mathrm{SNR}_{\mathrm{bin}} = P_{\mathrm{sig}}/P_{\mathrm{noise}}$ and reported in dB.

Fig. 5 shows the experimental BER. It compares spectrogram based decoding at OVRO with experimental results from [10] which used raw IQ samples as input, indicating minimal performance loss. The error probability drops below 10^{-2} when SNR is higher than about -7 dB and approaches 10^{-4} around -6 dB. This sharp transition happens because correlation peaks become stable and chip averages separate from noise.

At very low SNR (< -10 dB), correlation peaks become less significant and chip energy separation is weak, so decod-

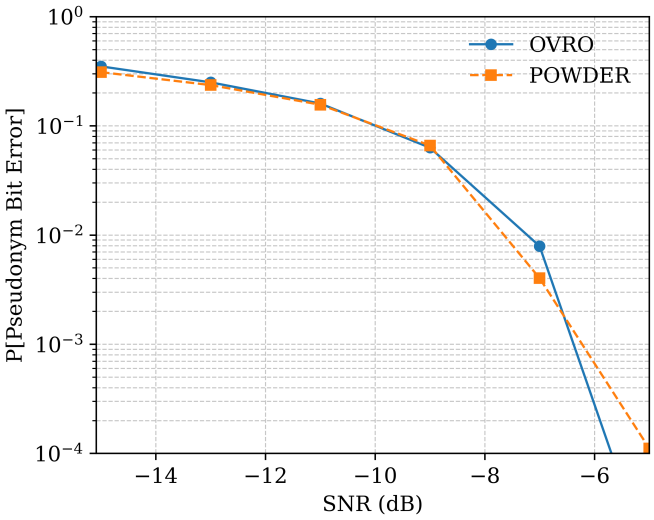

Fig. 5: Probability of pseudonym bit error vs. SNR, for spectrogram-based OVRO backend experiment and IQ-based baseline results from POWDER [10]. Each OVRO point represents results over 5040 pseudonym bits.

ing errors increase. Forward error correction (FEC) was not evaluated in this study; incorporating FEC would be expected to further reduce pseudonym BER.

These results demonstrate that passive receivers can reliably detect and identify interfering devices from spectrogram data outputs. This capability enables accountability-based reporting and supports closed-loop interference management in shared-spectrum environments.

VI. Conclusion and Future Work

This work validates the practical feasibility of pseudonym watermark decoding using only spectrogram outputs from an operational radio astronomy backend, without modifying telescope hardware or accessing IQ data. RFI detection and identification performance depends on the backend resolution and watermark resource allocation rather than access to baseband samples. In the present implementation, decoding assumes a single active interfering transmitter at a time. The watermark bandwidth may span one or more spectrogram bins depending on backend frequency resolution Δf_{rx}. If multiple interfering devices transmit simultaneously, their watermark energy superposes within the same spectral region, reducing separability under energy-domain decoding. Extending the spectrogram-based receiver to robust simultaneous multi-user operation remains future work.

Performance follows the rate–integration trade-off derived in Section IV. Backend frequency resolution Δf_{rx} and time integration determine the maximum reliable pseudonym bit rate, while integration gain enables attribution at low per-bin SNR. Thus, spectrogram-based pseudonym decoding is backend-dependent. Passive receivers with different receive parameter configurations may operate under different attribution capacity and rate constraints.

Prior work on pseudonymetry [10], [14] assumes cooperative transmitters embedding identifiable watermarks in their waveform, enabling device-level identification at the receiver. Non-cooperative interferers can still be detected but cannot be attributed to a specific identity. When integrated into coordinated spectrum-sharing frameworks, this mechanism provides device-level accountability using standard backend spectrogram outputs. These results demonstrate that pseudonymetry, validated on an operational radio astronomy backend, can support more accountable spectrum sharing between active and passive systems. This work provides a technical foundation for more cooperative and efficient coexistence in future spectrum environments.

Acknowledgements

This material is based upon work supported by the US National Science Foundation under Grants #2229427, #2346555, and #2535003.

References

[1] C. Wilson, "Propagation prediction in establishing a radio quiet zone for radio astronomy," in *European Conference on Antennas and Propagation (EuCAP)*, pp. 1209–1213, 2014.

[2] S. W. Ellingson, "Radio astronomy interference mitigation: Current status and future directions," in *IEEE URSI General Assembly*, 2005.

[3] J. Ford, C. Anderson, and G. Hampson, "Rfi mitigation and signal processing in radio astronomy," *Publications of the Astronomical Society of Australia*, vol. 31, 2014.

[4] A. R. Offringa, G. Bernardi, and G. Ghirardini, "The low-frequency array (lofar) rfi mitigation strategy and results," *Astronomy & Astrophysics*, vol. 574, p. A61, 2015.

[5] A. Anastasopoulos, C. K. Anagnostopoulos, and A. G. Kanatas, "Interference mitigation in satellite services coexisting with passive services," in *IEEE International Symposium on Wireless Communication Systems (ISWCS)*, pp. 1–5, 2013.

[6] N. Papadopoulos, M. Lofquist, A. W. Clegg, and K. Gifford, "Spectrum sharing dynamic protection area neighborhoods for radio astronomy," in *IEEE Wireless Communications and Networking Conference (WCNC)*, pp. 1–6, 2023.

[7] S. Munira, D. Saha, G. Hellbourg, and A. Dutta, "Dynamic protection zone for radio astronomy," in *Proceedings of IEEE/ACM Dynamic Spectrum Access Networks (DySPAN)*, 2024.

[8] Y. Zhang, F. Yan, K. Gong, and J. Ma, "Deep learning approaches for radio frequency interference detection in radio astronomy data," *Monthly Notices of the Royal Astronomical Society*, vol. 514, no. 2, pp. 2345–2357, 2022.

[9] M. G. Weldegebriel, J. Wang, N. Zhang, and N. Patwari, "Pseudonymetry: Precise, private closed-loop control for spectrum reuse with passive receivers," in *2022 IEEE International Conference on RFID (IEEE RFID)*, pp. 91–96, 2022.

[10] M. Weldegebriel, Z. Li, D. Maas, G. Hellbourg, N. Zhang, and N. Patwari, "Sensing and stopping interfering secondary users: Validation of an efficient spectrum sharing system." arXiv preprint, 2025.

[11] M. B. Sherman, N. Kosogorov, C. Law, V. Ravi, J. T. Faber, S. K. Ocker, L. Connor, Y. Qu, K. Shin, K. Sharma, *et al.*, "Deep synoptic array science: Searching for long duration radio transients with the dsa-110," *Publications of the Astronomical Society of the Pacific*, vol. 138, no. 2, p. 024501, 2026.

[12] V. Ravi and DSA-110 Collaboration, "The DSA-110: overview and first results," in *American Astronomical Society Meeting Abstracts*, vol. 241, pp. 239–01, 2023.

[13] International Telecommunication Union, "Protection criteria used for radio astronomical measurements (recommendation itu-r ra.769-2)," *ITU Recommendations*, 2003.

[14] M. G. Weldegebriel, J. Wang, G. Hellbourg, N. Zhang, and N. Patwari, "Watermarking of OFDM for pseudonymetry: Analysis and experimental results," in *Proceedings of the IEEE International Conference on Communications Workshops (ICC Workshops)*, pp. 317–322, 2024.