# Sensing and Stopping Interfering Secondary Users: Validation of an Efficient Spectrum Sharing System

Meles G. Weldegebriel
Electrical and Systems Engineering
Washington University in St. Louis
St. Louis, MO, USA

Zihan Li
Computer Science and Engineering
Washington University in St. Louis
St. Louis, MO, USA

Dustin Maas
Kahlert School of Computing
University of Utah
Salt Lake City, UT, USA

Gregory Hellbourg
Cahill Center for Astronomy & Astrophysics
California Institute of Technology
Pasadena, CA, USA

Ning Zhang
Computer Science and Engineering
Washington University in St. Louis
St. Louis, MO, USA

Neal Patwari
Price College of Engineering
University of Utah
Salt Lake City, UT, USA

***Abstract*—We present *Stoppable Secondary Use (StopSec)*, an interference management protocol for spectrum sharing systems that enables primary users (PUs) to identify and quickly halt interference from secondary users (SUs). StopSec embeds a lightweight signal-level watermark into SU transmissions, allowing PUs to robustly recover a one-time pseudonym even under low-SNR and time-varying channel conditions, and record interference events in a shared database without exposing SU identities. Each SU periodically checks this database and vacates the channel if their pseudonym appears, enabling precise, privacy-preserving interference control without relying on persistent identifiers or conservative exclusion zones. We introduce a single-subcarrier watermark for OFDM-based SU links and a coded modulation robust to practical time-varying channels. We implement and evaluate StopSec through extensive real-time over-the-air experiments. We validate that StopSec does not degrade the SU data link. Interfering SUs can be stopped in under 150 ms, even under real-world channel variations. Even when the SU signal is 10 dB below the noise floor, and when multiple SUs simultaneously interfere with the PU, StopSec successfully identifies and stops interference within seconds. These results demonstrate that StopSec provides an effective, scalable foundation for automatic and accountable spectrum sharing.**

***Index Terms*—Dynamic Spectrum Sharing, Interference Mitigation, Signal Watermarking**

## I. Introduction

Highly dynamic spectrum sharing is essential to maximizing the benefits of limited spectrum resources across diverse users [1], [2]. Today's sharing systems are conservatively designed for the *worst-case* interference, to ensure secondary users (SUs) never interfere with primary users (PUs), but this results in overly large exclusion zones. For example, the conservative propagation model used for the Citizens Broadband Radio Service (CBRS) band resulted in dynamic protection areas (DPAs) being about 5× larger than necessary to protect US naval radars [2]. We see the push for conservative design from incumbents also in the case of secondary unlicensed use of the 6 GHz band through an Automated Frequency Coordinator (AFC). In this case, incumbent AT&T sued [3], arguing that despite the conservative propagation model [4], interference "inevitably" occurs, and the lack of a mitigation mechanism made their microwave backhaul systems vulnerable. Indeed, if we have fast and automated reactive mitigation mechanisms to recover from (and remove) interfering SUs, we can avoid the pitched battle between efficient spectrum sharing and protected incumbents, and achieve both.

However, we have yet to see a reactive interference management system able to quickly stop an SU from operating on the band if it is the device interfering with a protected PU, without negatively impacting the SU communication link. Manual enforcement is too slow — it can take years to disable rogue transmitters [5]. One existing reactive approach is to turn off subsets of SUs until the interference is no longer observed [6], but this can be slow, and some SUs are turned off despite not interfering. One could identify interfering SUs using identifying data in a SU's packet header, but PUs can be interfered with at power levels many dB below that required for header demodulation. An extension is for SUs to watermark their RF transmissions to allow primary receivers to detect and identify the specific interfering user [7], [8]. However, existing RF watermarking methods (1) can degrade the secondary wireless link by 3 dB [7], (2) can add several ms of timing jitter [8], which is not acceptable for applications like 5G low latency, and (3) can fail in time-varying mobile wireless channels. To date, no published work has implemented RF watermarking over-the-air (OTA) to stop interference in real time.

We propose, implement, and experimentally validate *Stoppable Secondary Use* (StopSec), a protocol that enables real-time detection and stopping of the particular SU whose transmission is interfering with a PU in a spectrum sharing system. We assume a future standardized highly dynamic spectrum sharing protocol followed by SUs and used by PUs to coordinate the closed loop control. SUs in StopSec pick random pseudonyms and watermark their transmissions. If they interfere with a PU, even down to 10 dB below the noise floor, the PU decodes the pseudonyms, and logs them in a shared, time-aware database. SUs periodically query the database and switch out of the channel if any of their pseudonyms have been reported. Our framework supports low-overhead coordination, minimizes false positives, and preserves SU privacy through unlinkable pseudonym generation.

Wireless channels are time-varying, and low-rate watermarking methods such as [9] are sensitive to changes in the channel, and in the noise and interference levels. As we intend to detect interference that is on the order of 10 dB below the noise floor, even small channel changes are problematic. We solve this by using a time-varying code in the pulse shape of the watermark modulation that allows a receiver to cancel channel changes and focus on the watermarked bits. Further, the modulation is designed to make reception simple — the watermark receiver uses energy detection, and no phase or frequency synchronization is required. This is a critical requirement for reliable operation when the interfering signal-to-noise ratio (SNR) is very low.

Further, we address the problem that a watermark can degrade an SU communication link. In StopSec, which is designed for multicarrier modulated signals, we chose to encode pseudonym bits in just one of the subcarriers, in contrast to modifying the entire SU communication signal as in [7], [8]. The impact is comparable to a pilot channel, and is $\leq 1.5\%$ of the bandwidth in typical OFDM systems.

We make the following contributions in this paper:

- We introduce the design for StopSec. This includes a new watermarking scheme, Code Pulse Amplitude Modulation (CPAM), which encodes each pseudonym bit as a structured received power pattern distributed across multiple signal intervals. CPAM provides resilience to fading, power fluctuations and time-varying channels. StopSec applies CPAM to a single subcarrier of an OFDM signal, minimizing impact on the data communication on the SU link, and enabling stopping of multiple SUs that interfere with the PU simultaneously.
- We implement StopSec and use it to perform extensive real-time full system OTA evaluation on the POWDER wireless testbed. Open source implementation code is at [10]. StopSec is demonstrated to achieve sub-second interference suppression, even at $-10$ dB SNR or with overlapping interfering SUs.

## II. System Design

### A. System Overview

StopSec operates a spectrum-sharing system with three major entities: secondary users (SUs), primary users (PUs), and an interference-reporting database, as shown in Figure 1.

At a high level, StopSec defines an RF watermarking method that allows PUs, if they experience interference, to identify and notify the SUs responsible for causing the interference. The watermark is designed to remain robust under fluctuating channel conditions, be detected reliably at very low received power, and consume minimal bandwidth. StopSec watermarks with one-time ***pseudonyms*** to ensure SU privacy, as these pseudonyms are not tied to any persistent SU identity.

More specifically, we consider SUs which use an OFDM-based communication system that shares the spectrum with the PU but operates with lower priority. To enable StopSec, each SU embeds a self-generated random pseudonym into its OFDM transmissions and records the pseudonym along with the time and channel of use. The SU then periodically checks the database to determine whether its transmissions have interfered with any PU. If a pseudonym in the database matches any record of the SU, the SU will vacate the affected channel to avoid causing interference to PUs.

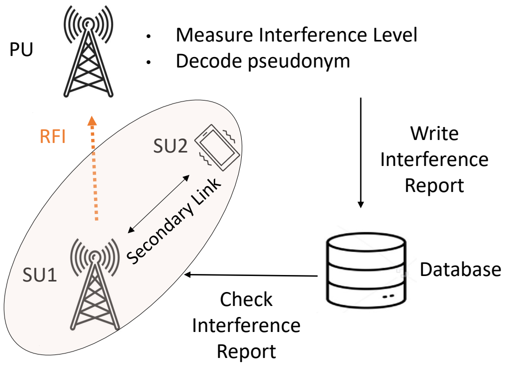
PU
Measure Interference Level
Decode pseudonym
RFI
SU2
Secondary Link
SU1
Write Interference Report
Database
Check Interference Report

Fig. 1. StopSec system model.

For the PU, StopSec requires that when interference is detected, the PU recovers the pseudonym embedded in the interfering signal and logs an *interference report* in the database. This report includes the recovered pseudonym, timestamp, location, and channel, and serves as a notification to the corresponding SU to take action.

The last entity, the database, stores recent interference reports from PUs and provides the feedback loop that informs an SU when it is causing interference. In our design, the PU writes interference reports, while the SU performs only read operations to obtain the information needed to adjust its transmission parameters (e.g., switching channels).

In the following sections, we provide detailed descriptions of how each component of StopSec is designed. We first present the RF signal watermarking design, then describe the pseudonym and packet-level design, and finally explain how the pseudonym functions within the full control loop.

### B. Signal Watermarking

At the signal level, there are major challenges to address. The first challenge is that SU signals can interfere with a PU's normal operation even if the SU signals are below the PU's noise floor. We enable watermark detection at low SNR using a low rate watermark designed for non-coherent reception. The second challenge is fluctuating wireless channels and other (non-watermarked) interference sources which cause time-variation in SU signal and noise powers at the PU that make watermark decoding difficult. We address this by using a *coded pulse amplitude modulation (CPAM)* method. The third challenge is ensuring that StopSec does not interfere with the SU's normal communication process. We address this by realizing StopSec within *a single subcarrier*.

*1) Watermark for Low SNR Detection:* In practical spectrum-sharing systems, PUs need protection from interference that is even below its noise floor. For example, the 6 GHz AFC requires interference to remain at $-6$dB relative

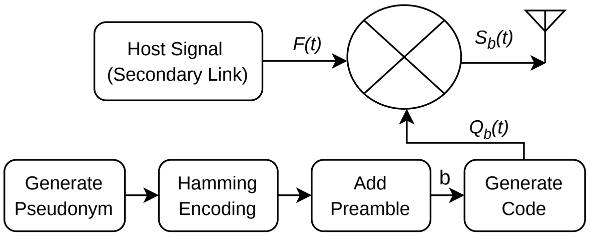

Fig. 2. Framework for CPAM watermarking.

to noise [3], [11]. At such low SNRs, an SU's packet data cannot be reliably demodulated. Further, phase and frequency synchronization, which are prerequisites for demodulation of phase-shift keying (PSK) and quadrature amplitude modulation (QAM) methods, is difficult at low SNR. Finally, we want the pseudonym receiver at the PU to be low complexity, to minimize the cost for PUs.

StopSec uses amplitude modulation, and detection using the average power of the received signal, which is referred to as energy detection. This enables demodulation without phase and frequency synchronization, in fact, without direct use of IQ samples. By using a low bit rate for the watermark, we design StopSec to be detectable at a very low SNR.

*2) Coded Pulse Amplitude Modulation (CPAM):* While energy detection enables low SNR detection, it can be confused by other changes in the channel. In mobile wireless channels, both the signal (the watermark from the SU) and the ambient level of noise and (non-watermarked) interference changes rapidly over time. If using PAM, in which each bit $b \in \{0, 1\}$ is mapped to a single amplitude, channel changes can overwhelm the constant amplitude. In contrast, CPAM encodes each bit as a sequence of $L$ changing (but known) amplitudes $\{A_{b,0}, \ldots, A_{b,L-1}\}$, spread across $L$ chips within the symbol period. Formally, to create the CPAM signal, we construct one pseudonym symbol $Q_b(t)$ as

$$Q_b(t) = \sum_{l=0}^{L-1} \left[1 - A_{b,l}\alpha\right] \phi(t - lT_c), \tag{1}$$

where $l$ indexes the chip within the symbol, $\alpha$ is the modulation index, $T_c$ is the chip duration, and $\phi(t)$ denotes the chip pulse shape (a rect function in OFDM). Since CPAM uses a changing amplitude pattern, a receiver can separate the time-varying fading from the code, and perform reliable detection even at extremely low SNRs. Since all of our experiments transmit one bit per symbol, we use the terms *pseudonym bit* (or *p-bit*) and *pseudonym symbol* interchangeably in this paper.

**Comparison to DS-SS**: A CPAM receiver correlates the received power over the $L$-chip sequence of known power levels, $\{|A_{b,0}|^2, \ldots, |A_{b,L-1}|^2\}$. CPAM offers substantially improved robustness to noise and channel fluctuations in a similar manner to direct-sequence spread spectrum (DS-SS), but CPAM operates on the signal power, not its phase. Further, rather than assigning unique orthogonal codes to users like DS-SS, CPAM uses a single code for all SUs to reduce receiver complexity. We show in Section IV-D that StopSec can still detect and stop multiple simultaneously interfering SUs.

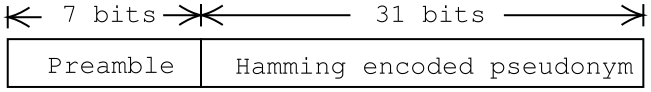

Fig. 3. Frame structure for pseudonym packets.

As shown in Figure 2, the CPAM code $Q_b(t)$ modulates the pseudonym subcarrier $f(t)$ to produce the watermarked signal $S_b(t)$ via $S_b(t) = Q_b(t)\, f(t)$. In the next section, we examine the case where $f(t)$ corresponds to a single subcarrier of the SU's OFDM signal over the entire OFDM packet duration. Because the watermark is low rate, a single OFDM packet might carry only one pseudonym bit, in which case a complete pseudonym is embedded across multiple OFDM packets. To isolate the effect of CPAM, we also evaluate CPAM separately from single-subcarrier watermarking, in which case $f(t)$ represents the entire SU data packet, in Section IV-A1.

*3) Single Subcarrier Watermarking:* The second challenge is to ensure that StopSec does not degrade the SU's normal communication link performance. Existing watermarking methods often introduce measurable penalties to the data link, such as reduced demodulation performance or additional timing jitter [9], [8], which may be unacceptable in cooperative spectrum-sharing scenarios. In contrast, StopSec is designed to impose minimal or no impact on the SU's data channel.

To achieve this, we reserve a single OFDM subcarrier as a dedicated watermarking subcarrier, leaving all other subcarriers untouched. In typical multi-carrier systems such as IEEE 802.11ax, dedicating one subcarrier (e.g., 1 out of 256) represents negligible overhead and is comparable to existing pilot or guard-band subcarriers used for link reliability. Moreover, because the watermarked subcarrier changes slowly relative to the OFDM data symbol rate, it has low spectral sidelobes and can naturally serve in place of a guard-band subcarrier. For reliable low-SNR detection, we use full modulation, i.e., $\alpha = 1$ in (1), meaning the pseudonym subcarrier is either silent or transmitted at twice the amplitude of a data subcarrier. In our experiments, the watermark is embedded either in the DC subcarrier or one of the guard-band subcarriers. The study of optimal subcarrier selection for watermarking remains an important direction for future work.

### *C. Pseudonym Generation and Detection*

StopSec uses a frame structure for each pseudonym packet consisting of two fields: a 7-bit preamble and a coded pseudonym, as shown in Figure 3. The preamble is a maximal-length sequence (m-sequence) [12], selected for its strong autocorrelation properties. It marks the start of the frame and enables the receiver to synchronize and filter out corrupted or spurious transmissions.

The coded pseudonym follows the preamble and represents the SU's temporary identity in a privacy-preserving manner. Each SU generates a fresh random pseudonym with no device-specific information. Each pseudonym comprises 26 random

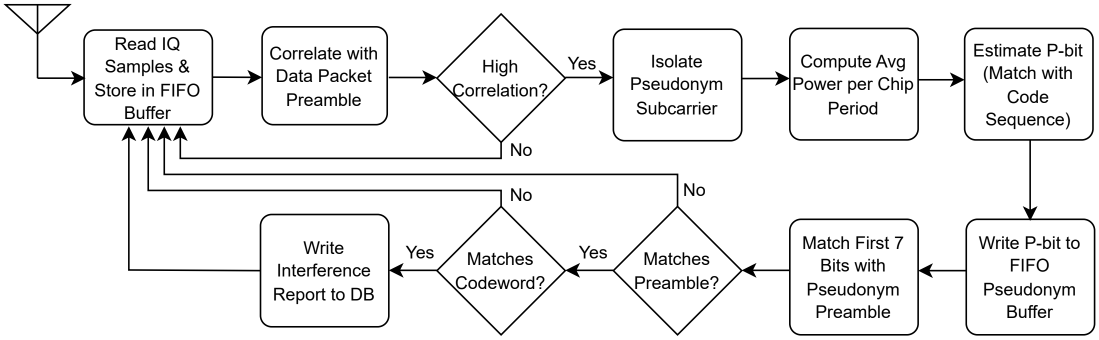
Read IQ Samples & Store in FIFO Buffer
Correlate with Data Packet Preamble
High Correlation?
Yes
No
Isolate Pseudonym Subcarrier
Compute Avg Power per Chip Period
Estimate P-bit (Match with Code Sequence)
Write P-bit to FIFO Pseudonym Buffer
Match First 7 Bits with Pseudonym Preamble
Matches Preamble?
Yes
No
Matches Codeword?
Yes
No
Write Interference Report to DB

Fig. 4. Pseudonym detection and reporting flowchart. The PU processes IQ samples to detect interference, demodulate and decode pseudonym bits, and log interference reports in real time.

bits, encoded using a (31, 26) Hamming code [13]. This forward error-correcting code supports single-bit correction and double-bit error detection, adding robustness to pseudonym detection capability.

The use of a 26-bit pseudonym has a balance between scalability and reliability. Specifically, it limits the probability of false positives during suppression. If a pseudonym consists of $K$ random bits and only one SU is transmitting, the probability that a second SU coincidentally uses the same pseudonym — at the same time, on the same channel is $2^{-K}$. In our case, this corresponds to a false alarm probability of $2^{-26} \approx 1.5 \times 10^{-8}$. In the worst case, a non-interfering SU is required to change channel due to a false alarm, but at this low probability, it is a negligible cost.

At the PU, pseudonym packets are accepted only if they pass two integrity checks:

1) **Preamble Correlation:** The receiver verifies that the initial 7 bits match the expected m-sequence. Frames failing this check are discarded.
2) **Error Control Coding:**. Hamming decoding corrects and validates the encoded pseudonym bits. Only validated frames lead to writing to the interference database.

These layered validation steps ensure high pseudonym detection accuracy, even under low-SNR conditions or when multiple SUs transmit concurrently.

At the PU side, pseudonym detection and reporting follow three steps (as shown in Figure 4):

1) **Interference Detection via Adaptive Thresholding:** The PU continuously correlates incoming samples with the OFDM packet preamble and applies an adaptive peak-to-median thresholding rule—similar to CFAR detection [14]—to distinguish OFDM signals from noise. Detection is declared when *peak_correlation* $\geq$ *threshold_factor*$\times$*median_correlation*, where the threshold factor begins at an empirically chosen value (4.5) and adjusts dynamically based on recent detection outcomes. This adaptive strategy enables StopSec to operate reliably with a fluctuating noise floor and at low SNR.
2) **Pseudonym Extraction and Validation:** Once a signal is detected, the PU isolates the pseudonym subcarrier using an FFT or narrowband filter and estimates each p-bit by correlating chip-level power values with the known CPAM code. The recovered p-bits are assembled into a pseudonym frame and validated using a 7-bit m-sequence preamble, followed by a $(31, 26)$ Hamming decoder to correct single-bit errors and detect double-bit errors. Only frames that pass these checks are accepted as valid pseudonyms.
3) **Interference Reporting:** After decoding a valid pseudonym, the PU records it with a timestamp and channel and writes an interference report to the shared database. This database serves as the feedback channel for SUs. When an SU finds its pseudonym in the database at a time and channel it used that pseudonym, it infers that it caused interference and vacates that channel. In our single-PU, single-channel experiments, only the pseudonym and timestamp are required.

### D. Database and Scalability

As the number of SUs increases, the database must remain efficient in both storage and query latency. Let $U_s$ denote the number of SUs, $R$ the average number of reports per user, and $S_r$ the report size. Since StopSec is designed for scenarios where interference is rare, $R \ll 1$, and the database size is $D = U_s \times R \times S_r = O(U_s)$, indicating linear growth in the number of SUs.

We implement the interference database using SQLite [15], a lightweight embedded system well suited for frequent low-overhead read operations. A time-to-live expiration mechanism removes stale entries after a fixed time-to-live $T_0$, and matched pseudonyms are deleted immediately, preventing unnecessary growth and reducing false positives. To support larger deployments, multiple localized databases can be used to limit geographic scope and reduce load.

Query latency depends on the number of concurrent users. If $T_b$ is the baseline response time and $U_Q$ is the number of concurrent queries, the expected response time can be approximated by $T_{\text{resp}} = T_b(1 + \omega(U_Q - 1)) = O(U_Q)$, where $\omega$ is the constant factor for the computation overhead. In practice, SQLite handles many concurrent reads efficiently, while writes incur higher overhead. This asymmetry matches our usage model: SUs perform frequent reads, whereas PUs write only when interference is detected. Prior work [16] also shows that SQLite performs competitively under concurrent workloads, further validating its suitability for StopSec.

### *E. Security and Privacy*

One of the key privacy features of StopSec is its use of pseudonyms, which are random bit strings with no link to device identity. Because they are freshly generated and independent of the data payload, pseudonyms do not reveal transmitter identity or increase the information exposed to nearby observers. However, privacy threats can still arise. An eavesdropper may demodulate a pseudonym and launch a man-in-the-middle attack by relaying it to a malicious transmitter near the PU. This could cause the pseudonym to appear in the interference database and unjustly force the legitimate SU to vacate the channel. Such attacks can be mitigated by checking timestamp consistency in interference reports or using source-localization techniques to identify active adversarial transmitters.

The database also introduces potential vulnerabilities. An attacker could inject fake pseudonyms or attempt to disable the database. To protect integrity, StopSec employs role-based access control, allowing only authorized PUs to write entries, while SUs have read-only access. Secure, authenticated APIs further prevent unauthorized modifications and provide traceability. To guard against denial-of-service attacks, redundant or distributed database deployments can improve availability.

## III. Implementation

### *A. Experimental Setup*

We deployed our experimental setup on the POWDER wireless testbed [17] using four rooftop base stations. Each base station is equipped with a National Instruments (NI) USRP X310 SDR and an associated compute node for transmission and reception processing. Three of the base stations serve as SUs, while the fourth functions as the PU. As illustrated in Figure 5, SU1, SU2, and SU3 are located approximately 330m, 590m and 570m from the PU, respectively. All users operate at a 3.385 GHz center frequency within the lower part of the cellular C-band. Depending on the experiment, SUs transmit OFDM waveforms on channel bandwidths ranging from 1 MHz to 10 MHz, spanning from 64 to 256 subcarriers. All system components — including watermark embedding, pseudonym detection, interference database access, and SU control — were developed in Python using UHD APIs for low-level radio control and signal processing. This setup enables end-to-end evaluation of StopSec's performance in a large-scale real-world outdoor deployment.

### *B. Pseudonym Watermarking Implementation*

To explore the trade-offs between detection reliability and data integrity, we implement two watermarking schemes: (i) a full-band CPAM approach that distributes watermark energy across all subcarriers, and (ii) a single-subcarrier CPAM scheme, which we refer to as the StopSec approach, that isolates pseudonym transmission to a dedicated subcarrier. This allows us to experimentally quantify the benefits and drawbacks of both strategies that we propose in StopSec, i.e., coded pulse amplitude modulation, and single subcarrier.

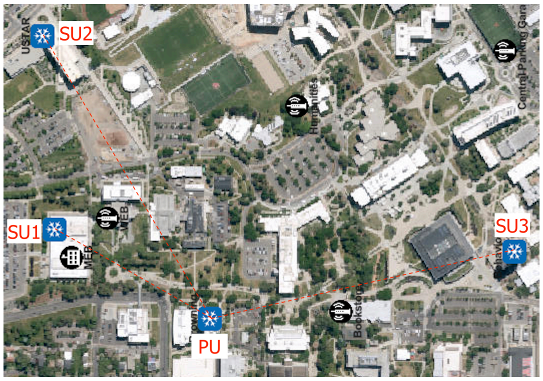

Fig. 5. Experiment map: SUs (SU1, SU2 and SU3) all interfere with the PU (bottom middle).

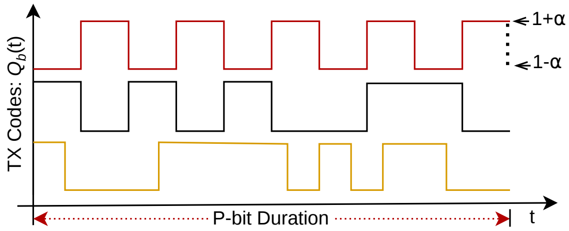

Fig. 6. Codes do not need to be m-sequences. Top: an alternating chip amplitude pattern with a chip length of 10. Middle: a truncated m-sequence, chip length 10. Bottom: a full-length m-sequence, chip length 15.

The code patterns in Figure 6 are representative variants of our watermarking mechanism in (1). We selected these examples to evaluate different code implementations with ample chip-to-chip transitions and balanced amplitude levels, allowing us to compare their pseudonym detection performance.

In our experiment, each SU embeds one pseudonym bit per data packet. Each OFDM packet contains 100 symbols, with the number of subcarriers determined by the transmission bandwidth. For example, a 2 MHz configuration uses 64 subcarriers: 48 data, 4 pilots, and 12 guard bands, with all data subcarriers modulated using QPSK. We configure the OFDM packets to begin with the high-throughput short training field (HTSTF) symbols defined in 802.11n [18], which we leave unwatermarked to preserve synchronization performance.

### *C. Pseudonym Reception and Reporting*

We implemented pseudonym detection and reporting entirely in Python, using the UHD Python API for interfacing with USRP X310 SDRs—–prioritizing accessibility over execution speed optimization. Still, this implementation runs in real time on a Dell R430 server equipped with two Intel E5-2630v3 8-core CPUs at 2.4 GHz.

### *D. System Configurations for Evaluation*

To assess StopSec's performance under diverse operating conditions, we varied key PHY-layer parameters — including subcarrier counts, bandwidths, and the number of interfering SUs — as summarized in Table I.

TABLE I
SYSTEM CONFIGURATIONS IN STOPSEC EVALUATION.

| Bandwidth | Subcarriers | Subcarrier Spacing | Pseudonym Bandwidth | Concurrent SUs |
|---|---|---|---|---|
| 2 MHz | 64 | 31.25 kHz | 1 / 64 = 1.5% | 1–3 |
| 5 MHz | 128 | 39.06 kHz | 1 / 128 = 0.8% | 1–3 |
| 10 MHz | 256 | 39.06 kHz | 1 / 256 = 0.4% | 1–3 |

*1) Testing Across Multiple Subcarrier Settings:* To evaluate StopSec's sensitivity to subcarrier allocation, we implement three configurations under a fixed 2 MHz bandwidth: (i) a single pseudonym subcarrier (31.25 kHz), (ii) two subcarriers (62.5 kHz), and (iii) three subcarriers (93.75 kHz). When we use more than one subcarrier in StopSec, we simply apply the single-subcarrier coded pulse amplitude modulation to each subcarrier in the OFDM packet identically. All other transmission settings remain fixed. This setup helps assess the trade-off between detection robustness and spectrum overhead.

*2) Transmission Bandwidth and Subcarrier Configuration:* StopSec can be applied to OFDM signals of any bandwidth. We tested StopSec under three bandwidth settings: 2 MHz, 5 MHz and 10 MHz, corresponding to 64, 128, and 256 OFDM subcarriers, respectively. In all three configurations, a single subcarrier is dedicated to pseudonym watermarking, and the remaining subcarriers are allocated for data, pilots, and guard bands. The choice of bandwidth affects the subcarrier spacing, which is a critical factor for both watermark robustness and data performance. For example, at 2 MHz with 64 subcarriers, the subcarrier spacing is approximately 31.25 kHz, while at 5 MHz and 128 subcarriers, the spacing increases to 39.06 kHz. Wider subcarrier spacing improves frequency-domain separation, enhancing watermark detection at low SNR, though it may increase sensitivity to timing offsets at higher data rates.

*3) Scaling to Multiple Interfering SUs:* In the real world, sometimes, a second (or third) SU will start interfering with a PU before the system can shut off the first interfering SU. StopSec must be able to handle this case. To evaluate this scalability, we test StopSec with one, two, and three simultaneously active and interfering SUs. We use the experiment setup shown in Figure 5. Each SU transmits with a unique pseudonym but shares the same OFDM configuration. This setup enables us to analyze decoding latency, and detection reliability under increasing interference complexity. During multi-SU experimentation, all SUs started transmitting at the same time, and transmit OFDM packets continuously (without pausing between transmissions). Testing situations in which OFDM packets are spread out and interleaved with other users' packets over time remains important future work.

## IV. EVALUATION

In this section, we evaluate the performance of StopSec through a series of experiments. Our goal is to quantify how effectively the system can detect, identify, and suppress interfering SUs in real time under various operating conditions.

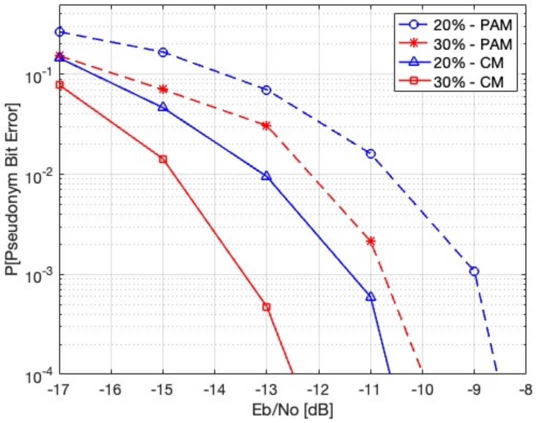

Fig. 7. Probability of pseudonym bit error vs. $\mathcal{E}_b/N_0$ at the PU receiver, comparing all-subcarrier CPAM (CM) and PAM watermarking schemes with modulation index $\alpha$ of 20% or 30%.

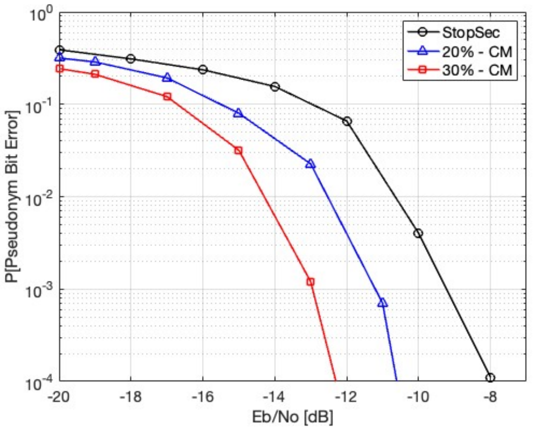

Fig. 8. Probability of pseudonym bit error vs. $\mathcal{E}_b/N_0$ at the PU. Plots compare all-subcarrier CPAM (CM, 20% and 30% modulation) vs. (single-subcarrier) StopSec.

### A. Performance of Watermarking Schemes

*1) Evaluation of Code Modulation:* We compare the pseudonym detection capabilities of all-subcarrier CPAM and PAM watermarking scheme of [9], [7]. Using the setup in Figure 5, we transmit all-subcarrier watermarked OFDM signals and measure the pseudonym bit error probability at the PU for each $\mathcal{E}_b/N_0$. We adaptively change the duration of the experiment following [19] — the number of experimental p-bits required to be experimentally transmitted is determined from the theoretical p-bit error rate in [7] and a desired 99% confidence level.

Figure 7 shows that all-subcarrier CPAM consistently outperforms PAM by more than 2 dB. For example, at $\mathcal{E}_b/N_0 = -11$ dB, CPAM achieves a pseudonym bit error probability of $7 \times 10^{-4}$, while PAM requires $\mathcal{E}_b/N_0 = -9$ dB to achieve similar performance.

*2) Impact of Watermarking on Data Demodulation:* To assess the impact of watermarking schemes on the performance of SUs, we evaluate the data bit error rate (BER) under three transmission configurations: (i) unwatermarked

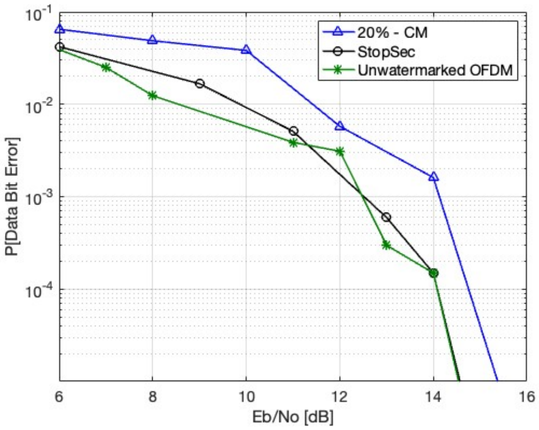

Fig. 9. SU OFDM receiver data BER vs. SU received $\mathcal{E}_b/N_0$, for $\alpha = 20\%$ all-subcarrier CPAM (CM), StopSec, and an unwatermarked OFDM signal.

OFDM (ii) StopSec watermarking (iii) 20% all-subcarrier CPAM watermarking. This experiment quantifies the effect of pseudonym embedding on data demodulation accuracy on the secondary link. In StopSec, we watermark only the dedicated pseudonym subcarrier using a modulation index $\alpha = 1.0$. We use $48$ data subcarriers, $4$ pilot subcarriers and $11$ guard subcarriers. For each configuration, we transmit OFDM packets over the secondary link and collect raw IQ samples at the receiver. The receiver performs standard OFDM baseband processing: symbol synchronization, channel equalization, 64-point FFT, and QPSK demapping to recover transmitted data bits. The experiment is repeated across a range of $\mathcal{E}_b/N_0$ values, which indicate the energy per data bit. To estimate $\mathcal{E}_b/N_0$, we measure noise power by sampling the channel (dedicated pseudonym subcarrier) when the transmitter is idle, and compute signal power from received packets during transmission. BER is computed by comparing received bits with the known transmit bits and calculating the ratio of bit errors to total bits received. To determine the number of data bits to send in order to reliably estimate the BER, we use the method of [19].

Experimental results shown in Figure 9 demonstrate that StopSec **does not** degrade data demodulation at the secondary receiver, whereas 20% all-subcarrier CPAM degrades performance by up to 2 dB. This validates that StopSec does not affect data demodulation at the secondary receiver.

### B. End-to-End Latency Evaluation

We evaluate the total latency from the onset of interference to the PU to the point at which the interfering SU is successfully stopped. This measurement captures the complete system response time, including: interference detection at the PU, pseudonym decoding from the watermarked subcarrier, write and lookup operations in the remote interference database, and delay until the SU query and stoppage of use of the band. Timestamps are recorded at each stage of the pipeline, enabling us to compute the total interference onset-to-suppression delay. This metric characterizes the responsiveness of the StopSec system under realistic deployment conditions.

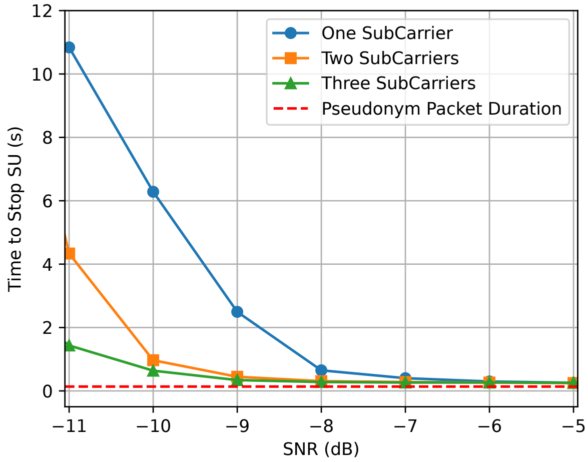

Fig. 10. Stopping latency vs. SNR. When StopSec can detect interfering SUs at lower SNR by using more (2-3) subcarriers.

Using the experimental setup illustrated in Figure 5, we evaluate the latency of the StopSec system. In this experiment, SU1 acts as the interfering SU, while PU serves as the PU monitoring the channel. We measure this latency across different watermarking subcarrier allocations and SNR levels. Figure 10 presents the results. The x-axis represents the SNR in dB, and the y-axis shows the latency (in seconds), i.e., the time it takes for StopSec to stop the SU from transmitting in the band.

As shown in Figure 10, the single-subcarrier configuration achieves an average latency of less than 270 ms when the SNR exceeds $-8$ dB. In contrast, both the two-subcarrier and three-subcarrier configurations maintain similar latency even at lower SNR levels ($-10$ dB), demonstrating an improvement in performance of approximately (2 dB). Further, the single-subcarrier configuration has latency of 150 ms at $-6$ dB SNR. Most of this latency is due to the duration of the pseudonym packet, which is 137 ms for these settings. These results indicate first that a single subcarrier StopSec system can reliably detect and stop interference caused by a SU within 270 ms, even when the interference is 8 dB below the noise power. Further, the results show how StopSec could be adapted to even lower SNR scenarios by adding additional subcarriers — 2-3 subcarriers allow detection with 2 dB less SNR.

### C. Impact of SU Transmission Bandwidth

To be incorporated into future spectrum sharing systems, StopSec must work with secondary systems that operate across a wide range of RF bandwidths. We contend that pseudonym detection and SU stopping performance is primarily determined by the pseudonym subcarrier bandwidth. To test this, we evaluated StopSec under three different transmission bandwidths - 2 MHz, 5 MHz, and 10 MHz - using 64, 128, and 256 subcarriers, respectively, in order to keep the subcarrier bandwidth relatively constant. In all configurations, a

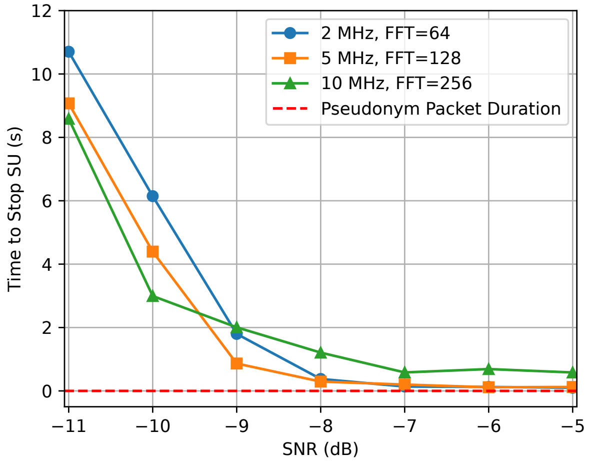

Fig. 11. Latency vs. SU signal bandwidth and number of OFDM subcarriers.

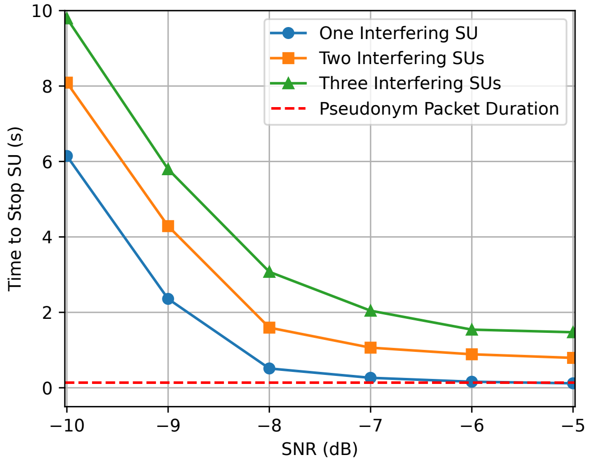

Fig. 12. Latency to stop all SUs vs. SNR and number of interfering SUs.

single subcarrier is allocated for pseudonym transmission, thus the pseudonym bandwidth overhead is approximately 1.5% (2 MHz), 0.7% (5 MHz), and 0.4% (10 MHz).

As shown in Figure 11, StopSec maintains robust interference detection and suppression performance across all bandwidth settings. Notably, the 5 MHz configuration demonstrates slightly superior performance compared to the 2 MHz case. This can be attributed to the somewhat larger subcarrier spacing in 5 MHz (39.06 kHz) compared to the 31.25 kHz subcarrier spacing in the 2 MHz setting.

However, in the 10 MHz configuration, we observe an increase in latency under high SNR conditions (above $-9$ dB). We determined that this degradation is due to the larger sampling rate required for processing 10 MHz transmissions, which increases the data throughput and computational complexity. At some rate, our real-time implementation of StopSec drops some samples, and thus pseudonym packets, before detection. This performance bottleneck is likely not a fundamental limitation, but an artifact of the low-compute implementation. It can be mitigated by deploying a more efficient detection algorithm or increasing processing resources at the primary receiver.

The results indicate that StopSec can be used across a range of SU signal bandwidths. Further, the required watermarking overhead decreases as a percentage of transmitted signal as the bandwidth increases, but StopSec performance remains approximately constant.

### *D. Stopping Multiple Interfering SUs*

*1) Evaluation of Latency Under Varying Number of Interfering SUs:* In real-world deployments, multiple SUs may interfere with a PU before the first can be shut down. To evaluate StopSec's scalability, we test scenarios with one, two, and three concurrently active SUs using the setup in Figure 5. Each SU randomly generates pseudonym packets while sharing identical OFDM settings. For each SNR value, we measure the latency under three conditions: (i) only SU1 transmitting, (ii) SU1 and SU2 transmitting, and (iii) all three SUs transmitting. In all cases, the PU runs the same detection pipeline—monitor the channel, decode pseudonyms upon detecting interference, and write the interference report to the database.

Figure 12 presents the latency results. As expected, latency increases as more interfering SUs become active. This is due to two factors: (i) pseudonym decoding becomes harder with overlapping interference, and (ii) the PU sequentially identifies and stops SUs one at a time. The single-SU case yields the lowest latency, followed by the two- and three-SU cases. Across all settings, latency is measured from the moment interference first arrives at the PU until all interfering SUs are shut down.

Despite the added complexity, latency remains below ten seconds for all three scenarios when the SNR is at least $-10$ dB, demonstrating the system's robustness and real-time responsiveness under moderately degraded conditions. During the experiment, SU1—whose signal is strongest at the PU—was consistently detected and mitigated first, followed by SU2 and then SU3. Interestingly, SU3 was stopped last despite being physically closer to the PU, owing to higher channel loss. This confirms that StopSec prioritizes SUs based on interference impact rather than physical proximity.

*2) Probability of Pseudonym Packet Detection:* Figure 13 presents the probability of successful pseudonym packet detection at the PU under two different interference scenarios: a single interfering SU and three concurrent interfering SUs. The y-axis denotes the probability of detection, while the x-axis represents the SNR at the PU in dB. As expected, the detection probability is consistently higher in the single-SU scenario compared to the three-SU scenario. This is due to increased interference in the presence of multiple concurrent transmissions, which degrades the PU's ability to reliably decode pseudonym packets.

In our context, a pseudonym packet is considered successfully detected only if it is transmitted by a SU, correctly decoded by the PU (including detection of the preamble and passing of the forward error correction check), logged in the

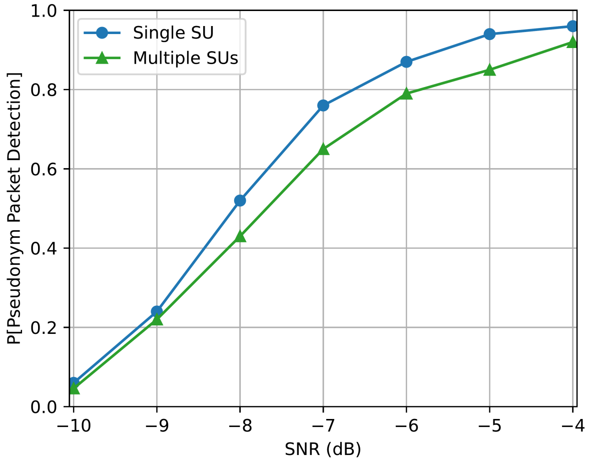

Fig. 13. Probability of pseudonym packet detection vs. SNR, comparing one interfering SU and 3 interfering SUs.

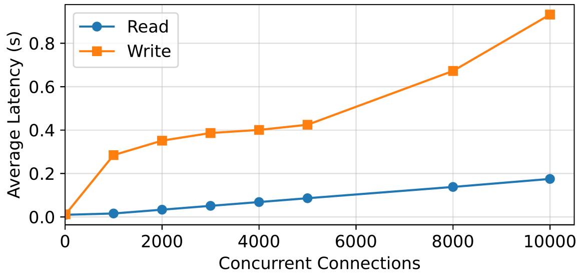

Fig. 14. Avg. database response time vs. concurrent connections.

remote database, and subsequently queried and matched with the local database, enabling timely suppression of the interfering SU. The observed degradation in detection probability under multiple interfering SUs highlights that it would likely take somewhat more time to suppress an interfering SU when multiple SUs are simultaneously interfering.

### *E. Database Scalability Evaluation*

To evaluate database scalability, we benchmark performance using Apache Bench [20], measuring response time under varying levels of concurrent requests. Figure 14 shows latency trends for both read and write operations as concurrency increases. As expected, write latency grows significantly because SQLite allows only a single writer at a time, forcing concurrent writes to serialize. In contrast, read operations scale efficiently: even with $10^4$ simultaneous read requests, the average response time remains below 200 ms. This behavior aligns well with our system workload, where PUs write only when interference is detected—an infrequent event—while SUs primarily issue read queries to monitor channel status. Consequently, the read-dominant access pattern in StopSec leverages SQLite's strengths, ensuring that the database remains scalable and does not become a performance bottleneck even as the number of SUs increases.

## V. Related Work

Maximizing radio spectrum efficiency and control is a priority for spectrum policy, radio science, and the wireless industry. There have been substantial advances in wireless technology and regulation, enabling more flexible use of shared bands. In this section, we summarize existing interference mitigation techniques used in spectrum sharing.

**RFI Detection and Suppression Schemes:** Traditional interference mitigation at passive receivers relies on receiver-side RFI detection and suppression, also known as "unilateral" mitigation [21]. Receivers filter, excise, or cancel to reduce interference [22], [23], [24], [25]. These approaches require no transmitter participation but degrade in high-occupancy environments and struggle with unknown or time-varying channels, especially at low SNR.

**Power Control Schemes:** Power control enables spatial reuse by keeping transmitted power below interference thresholds defined by the ITU [26], [27]. While effective, power control can be overly conservative when based on imperfect channel models and is typically combined with other sharing mechanisms.

**Cooperative Spectrum Sharing:** Cooperative sharing—involving coordination among incumbents and new entrants—is essential for meeting demand across commercial, scientific, and federal uses [28]. Coordination allows opportunistic access to incumbent bands [21], [29], and existing systems focus primarily on *ex ante* prevention [1].

However, any channel model will sometimes be far off. If we must ensure that interference only exceeds a threshold 5% of the time, we need 16 dB extra margin compared to the channel model's predicted received interference power, assuming the channel model of [30]. But if we can tolerate interference only 0.01% of the time, we need 37 dB margin. The additional 21 dB margin results in the area of exclusion being multiplied by a factor of 25, or equivalently, 96% of the possible sharing area being unused. Ensuring extremely rare interference leads to extremely inefficient spectrum sharing.

Future systems must mitigate harmful interference when it occurs. StopSec contributes to this direction by providing a closed-loop mechanism that detects, identifies, and suppresses interfering SUs in real time.

**Automated Frequency Coordination (AFC) and Spectrum Access System (SAS):** The AFC mechanism enables unlicensed operation in the 6 GHz band [31], [4], while the SAS manages the CBRS band (3550–3700 MHz) [32]. Neither sharing system can identify the specific interfering device when interference occurs [33], so both prevent interference to incumbents with location-based exclusion zones derived from conservative propagation models [34], [35], [36]. In the CBRS band, interference to incumbents has *never been reported* [37], which is said to indicate that models are far too conservative, reducing sharing efficiency [37], [38].

**Reactive Interference Control:** Reactive systems respond to real-time interference measurements. For example, the ASTRA framework [6] provides feedback to SUs based on measured aggregate interference power but cannot identify the specific interferer. StopSec complements this work by enabling per-transmitter accountability through passive PU sensing, RF watermarking, and distributed reporting.

## VI. Conclusion

StopSec provides a practical and scalable mechanism for dynamic spectrum sharing, enabling real-time detection and suppression of interfering SUs while requiring only a small fraction of SU bandwidth (e.g., 0.4% in a 10 MHz channel) and without degradation to SU communications. StopSec precisely targets only interfering SUs rather than relying on broad exclusion zones. Its low-rate CPAM watermark remains detectable even when SU interference is 10 dB below the PU's noise floor, even in time-varying channels, and supports suppression of multiple overlapping interferers. Our implementation on the POWDER testbed with one PU and three SUs demonstrates feasibility and aligns with emerging cloud-based coordination frameworks such as CBRS's SAS and the 6 GHz AFC. By enabling closed-loop feedback to shut off infrequent SU interference when it occurs, we can dramatically improve the efficiency of the next generation of spectrum sharing systems.

## Acknowledgements

The authors thank Shanghao Shi for his editing work. This material is based upon work supported by the US National Science Foundation under Grants #2229427, #2346555, and #2535003.

## References

[1] Wireless Spectrum R&D Interagency WG, NITRD, NSTC, "National spectrum research and development plan," Oct. 2024. National Science and Technology Council Working Group Report.

[2] A. Clegg, "Developments towards a more robust and dynamic spectrum sharing framework," tech. rep., WInnForum Working Document WINNF-TR-2016, July 2025.

[3] US FCC, "OET announces conditional approval for 6 GHz band AFC systems," 2021. Public Notice ET Docket No. 21-352.

[4] Y. Hassan, C. Carlos, A. Reza, and H. David, "Spectrum sharing using automated frequency coordination," Dec 2022. Intel White Paper.

[5] C. Welch, "Florida man drove around as a cellphone-jamming vigilante for two years," *The Verge*, May 2014.

[6] A. Sarbhai, D. Johnson, K. Webb, L. Stoller, O. Collaco, A. Orange, S. Zachary, B. Pearce, S. Tadik, S. Llosa, *et al.*, "Reactive interference management for radio astronomy in radio dynamic zones using ASTRA," in *2025 IEEE Intl. Symp. on Dynamic Spectrum Access Networks (DySPAN)*, pp. 1–10, 2025.

[7] M. G. Weldegebriel, J. Wang, G. Hellbourg, N. Zhang, and N. Patwari, "Watermarking of OFDM for Pseudonymetry: Analysis and Experimental Results," in *2024 IEEE Intl. Conf. on Communications Workshops (ICC Workshops)*, pp. 317–322, 2024.

[8] A. Palacios, D. Harman, C. Kitras, E. Kelsey, M. C. Burnett, W. K. Harrison, and P. Lundrigan, "Hidden in plain sight: Communicating using interference," in *2025 IEEE Intl. Symp. on Dynamic Spectrum Access Networks (DySPAN)*, May 2025.

[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 RFID*, pp. 91–96, 2022.

[10] M. G. Weldegebriel and N. Patwari, "Stopsec-protocol Github respository." https://github.com/Meles-Weldegebriel/StopSec-Protocol, 2026.

[11] S. Dogan-Tusha, A. Tusha, M. I. Rochman, H. Nasiri, J. R. Palathinkal, M. Atkins, and M. Ghosh, "Evaluation of indoor/outdoor sharing in the unlicensed 6 GHz band," *arXiv preprint arXiv:2505.18359*, 2025.

[12] S. W. Golomb, *Shift Register Sequences*. Aegean Park Press, 1982.

[13] S. Lin and D. J. Costello, *Error Control Coding: Fundamentals and Applications*. Prentice Hall, 2nd ed., 2004.

[14] F. Salahdine and N. Kaabouch, "Matched filter detection with dynamic threshold for cognitive radio networks," *arXiv preprint arXiv:1609.08398*, 2016.

[15] R. D. Hipp, "SQLite," 2020. https://www.sqlite.org/index.html, Accessed: 06.16.2025.

[16] K. P. Gaffney, M. Prammer, L. Brasfield, D. R. Hipp, D. Kennedy, and J. M. Patel, "Sqlite: past, present, and future," *Proceedings of the VLDB Endowment*, vol. 15, no. 12, 2022.

[17] J. Breen, A. Buffmire, J. Duerig, K. Dutt, E. Eide, M. Hibler, D. Johnson, S. K. Kasera, E. Lewis, D. Maas, A. Orange, N. Patwari, D. Reading, R. Ricci, D. Schurig, L. B. Stoller, J. Van der Merwe, K. Webb, and G. Wong, "POWDER: Platform for open wireless data-driven experimental research," in *14th Intl. Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH 2020)*, Sept 2020.

[18] IEEE, "IEEE standard for information technology—telecommunications and information exchange between systems—local and metropolitan area networks—specific requirements—part 11: Wireless LAN MAC and PHY specifications—amendment 5: Enhancements for higher throughput." IEEE Std 802.11n-2009, Oct. 2009. Available at: https://standards.ieee.org/standard/802_11n-2009.html.

[19] D. Mitić, A. Lebl, and Z. Markov, "Calculating the required number of bits in the function of confidence level and error probability estimation," *Serbian Journal of Electrical Engineering*, vol. 9, pp. 361–375, 2012.

[20] The Apache Software Foundation, "ab - Apache HTTP server benchmarking tool." https://httpd.apache.org/docs/current/programs/ab.html, Accessed: 06.16.2025.

[21] National Research Council, *Spectrum Management for Science in the 21st Century*. National Academies Press, 2010.

[22] C. Barnabaum and R. Bradley, "A new approach to interference excision in radio astronomy: Real-time adaptive cancellation," *The American Astronomical Society*, pp. 2598–2614, 1998.

[23] A. Leshem, A.-J. van der Veen, and E. Deprettere, "Detection and blanking of GSM interference in radio-astronomical observations," in *IEEE SPAWC Workshop*, pp. 374–377, 1999.

[24] M. E. Abdelgelil and H. Minn, "Non-linear interference cancellation for radio astronomy receivers with strong RFI," in *IEEE Global Communications Conf.*, pp. 1–6, 2017.

[25] M. Abdelgelil and H. Minn, "Impact of nonlinear RFI and countermeasure for radio astronomy receivers," *IEEE Access*, vol. 6, pp. 11424–11438, 2018.

[26] ITU-Report-2003, "Protection criteria used for radio astronomical measurements," *Recommendation RA.769-2*, March 2003.

[27] ITU-Report-2012, "Performance and interference criteria for satellite passive remote sensing," *Recommendation ITU-R RS.2017-0*, Aug. 2012.

[28] US FCC, "A preliminary view of spectrum bands in the 7.125–24 GHz range and a summary of spectrum sharing frameworks." Technical Advisory Council White Paper, Aug. 2023.

[29] H. Martikainen, M. Majamaa, and J. Puttonen, "Coordinated dynamic spectrum sharing between terrestrial and non-terrestrial networks in 5G and beyond," in *2023 IEEE Intl. Symp. on a World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pp. 419–424, 2023.

[30] ITU Radiocommunication Sector, "Guidelines for evaluation of radio transmission," tech. rep., ITU-R, Sep. 1997. M.1225.

[31] US FCC, "OET announces approval of seven 6 GHz band automated frequency coordination systems for commercial operation and seeks comment on C3 spectra's proposed AFC system," 2021. Public Notice ET Docket No. 21-352.

[32] Google Cloud, "Spectrum Access System." https://cloud.google.com/spectrum-access-system/docs/.

[33] J. Marsh, "AT&T statement on FCC order to allow unlicensed devices in 6 GHz band," April 2020.

[34] N. Hathiramani and R. Bernhardt, "Functional requirements for the US 6 GHz band under the control of an AFC system," tech. rep., WInnForum Document WINNF-TS-1014, April 2025.

[35] WInnForum Spectrum Sharing Committee, "Spectrum access system (SAS) to CBSD interface technical specification," winnf-ts-0016v1.2.7, Mar. 2022.

[36] WInnForum Spectrum Sharing Committee, "Cbrs technical requirements," tech. rep., 2018. Baseline Standards Release 1 for CBRS.

[37] A. Clegg, "Bridging the gap: Translating academic spectrum research into national impact," Sept 2025. Keynote presentation at NSF SWIFT / NewSpectrum PI Meeting.

[38] A. Sarbhai, F. Mitchell, S. Kasera, A. Bhaskara, J. Van der Merwe, and N. Patwari, "Reactive spectrum sharing with radio dynamic zones," in *2024 IEEE Intl. Symp. on Dynamic Spectrum Access Networks (DySPAN)*, pp. 429–438, 2024.