SweepSense: Sensing 5 GHz in 5 Milliseconds with Low-cost Radios

Yeswanth Guddeti, ***UC San Diego;*** **Raghav Subbaraman,** ***IIT Madras;***
Moein Khazraee, Aaron Schulman, and Dinesh Bharadia, ***UC San Diego***

https://www.usenix.org/conference/nsdi19/presentation/guddeti

This paper is included in the Proceedings of the 16th USENIX Symposium on Networked Systems Design and Implementation (NSDI '19).

February 26–28, 2019 • Boston, MA, USA

ISBN 978-1-931971-49-2

Open access to the Proceedings of the 16th USENIX Symposium on Networked Systems Design and Implementation (NSDI '19) is sponsored by

NetApp®

SweepSense: Sensing 5 GHz in 5 Milliseconds with Low-Cost Radios

Yeswanth Guddeti, Raghav Subbaraman[†], Moein Khazraee, Aaron Schulman, and Dinesh Bharadia
UC San Diego [†]*IIT Madras*

Abstract

Wireless transmissions occur intermittently across the entire spectrum. For example, WiFi and Bluetooth devices transmit frames across the 100 MHz-wide 2.4 GHz band, and LTE devices transmit frames between 700 MHz and 3.7 GHz). Today, only high-cost radios can sense across the spectrum with sufficient temporal resolution to observe these individual transmissions.

We present "SweepSense", a low-cost radio architecture that senses the entire spectrum with high-temporal resolution by rapidly sweeping across it. Sweeping introduces new challenges for spectrum sensing: SweepSense radios only capture a small number of distorted samples of transmissions. To overcome this challenge, we correct the distortion with self-generated calibration data, and classify the protocol that originated each transmission with only a fraction of the transmission's samples. We demonstrate that SweepSense can accurately identify four protocols transmitting simultaneously in the 2.4 GHz unlicensed band. We also demonstrate that it can simultaneously monitor the load of several LTE base stations operating in disjoint bands.

1 Introduction

High-time-resolution spectrum sensors [5, 18, 37, 32] enable new ways to share and manage the spectrum[1]. For example, the FCC granted permission for LTE providers to share licensed spectrum in the 3.5 GHz CBRS band with military radars, only if spectrum sensors are installed that can detect the military's millisecond-long military radar bursts anywhere within the 100 MHz bandwidth of the CBRS band [40]. In the future, we may even be able to improve co-existence of devices operating in the 5.8 GHz ISM band by performing high-time-resolution spectrum sensing of its 150 MHz bandwidth. For instance, a third-party high-time-resolution sensor can detect short intervals when WiFi devices are not using the spectrum, and inform unlicensed LTE base stations that they can operate without interfering [7].

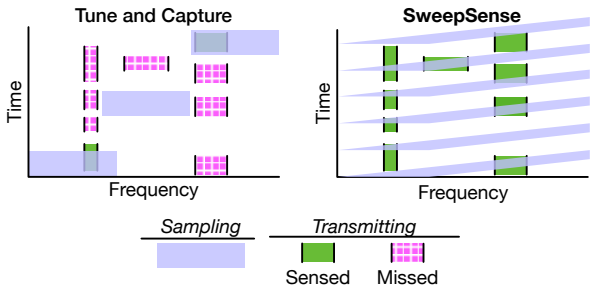

Figure 1: SweepSense rapidly sweeps its center frequency, rather than iteratively tuning and capturing the transmissions one frequency at a time.

Unfortunately, only complex and expensive spectrum sensors have both wide bandwidth and high time resolution. For example, there are radios that can sample several GHz of RF bandwidth continuously (e.g., OneRadio [14]). However, they are expensive (∼$500,000) due to their high-speed Analog-to-Digital converters, and complex due to the heavy computational power needed to perform real-time signal processing on high sample rates (e.g., GPUs or FPGAs). On the other end of the spectrum are narrow-bandwidth (∼50 MHz) radios (e.g., SDRs such as the USRP or HackRF [30, 37, 34]) that can not observe entire bands (e.g., 100 MHz) at once. The sensing bandwidth of these radios can be improved by intelligently tuning [37] but they are still likely to miss transmission due to their narrow bandwidth and the downtime they experience during tuning (as shown in Fig. 1).

We introduce a new paradigm in spectrum sensing, called SweepSense, which achieves both wide sensing bandwidth and high time resolution with off-the-shelf narrow-bandwidth radios. SweepSense introduces a fundamental shift in the receiver architecture of narrow-

[1]High-time-resolution spectrum sensors are defined by their capability to observe a portion of every transmission (e.g., packet).

bandwidth radios: instead of tuning into each frequency, sampling for a short time, then switching to the next frequency, SweepSense rapidly sweeps the frequency of the receiver across the spectrum (Figure 1). By sweeping rapidly across the spectrum, SweepSense achieves high time resolution with a narrow bandwidth radio. However, there are several challenges that we must overcome to demonstrate that SweepSense is practical and feasible:

Off-the-shelf radios do not sweep: SweepSense is only practical if it can be deployed on existing radios, such as SDR-based spectrum sensors [35]. The RF signal path on the radio should not require extensive changes to make it sweep. Prior approaches to making radios sweep by adding an expensive high-sample rate Digital-to-Analog converter that acts as a rapidly sweeping local oscillator are impractical [9].

Sweeping radios distort samples: Rapidly sweeping the center frequency of a radio results in samples that are collected at an unknown, and changing, center frequency. These samples need to be mapped to a single center frequency, and corrected for distortions introduced by sweeping. Furthermore, the continuous changing of frequency may reduce the sensitivity of the radio, making it impossible to detect weak signals.

Sweeping radios only visit bands for a short time: Rapidly sweeping radios collect a small number of samples in each band. This may break typical spectrum sensing-related signal analysis, such as signal type identification and spectrum occupancy detection.

We make the following contributions that address each of these challenges:

1. Making off-the-shelf radios sweep (Section 3): We show that with only a simple modification to the local oscillator circuit of a radio, we can make it rapidly sweep its center frequency. Specifically, we disconnect the feedback loop used to lock the receiver's local oscillator onto a specific frequency, and replace it with a sawtooth signal, thus making the center frequency sweep. We demonstrate the generality of this simple modification, by performing it on three of the most popular RF frontends for the USRP SDR, the WBX (50 MHz–2.2 GHz), SBX (400 MHz–4.4 GHz), and CBX (1.2 GHz–6 GHz).

2. Unsweeping samples (Section 4): We present a novel calibration and recovery process that corrects the continuously changing frequency in samples captured by the sweeping radio receiver. Specifically, we created a mechanism that inverts the effects of the sweeping center frequency by mixing it with complex conjugate of a calibration signal. Generating the calibration signal does not require any extra hardware: it is received through leakage from the radio's own RF transmitter (As TX loopback mode was not supported in the SDR). The result of the unsweeping process is a stream of samples that look as if they were collected at a fixed center frequency.

3. Evaluating analysis of short captures (Section 5): We demonstrate that even with the small number samples captured by SweepSense, the repeated patterns and unique features of the captured signals are retained. Specifically, we show that cyclo-stationary techniques when used in tandem with standard classification models need just 25 μsec captures of signals to classify accurately. Previously it was assumed that these techniques required capturing the entire transmission (e.g., $\sim$ 1 msec packet for WiFi).

We evaluate SweepSense by modifying a USRP N210 SDR to sweep, and performing experiments in both indoor and outdoor environments. We made the following observations: (1) SweepSense can classify signals with at least 90% accuracy (wideband DSSS and OFDM WiFi, as well as narrowband Zigbee and Bluetooth) with only 25 μs of samples, (2) SweepSense can simultaneously measure the millisecond level utilization of multiple LTE downlink channels over a bandwidth of 200 MHz. , and (3) SweepSense can accurately detect fleeting radar bursts, required for serving as a spectrum sensor for the CBRS spectrum.

The SweepSense implementation for the USRP N210 is open source and available at:
https://github.com/ucsdsysnet/sweepsense

2 Related Work

Spectrum sensing is an extensively studied area [23, 15, 42, 33, 41, 28, 24, 31]. Recent innovations have been focusing on improving the time resolution of spectrum sensors. To the best of our knowledge, SweepSense is the first work to suggest improving the time resolution of narrow-band spectrum sensors by making them rapidly sweep—without sacrificing their ability to classify transmitter type and characterize utilization. In this section, we describe how SweepSense complements, compares, and improves upon prior approaches to improving the time resolution of spectrum sensors.

HIGH-SPEED SPECTRUM ANALYZERS: The most common RF equipment that can sweep the spectrum quickly (i.e., tens of milliseconds) are high-speed spectrum analyzers, such as the Oscor Blue [32, 3, 36]. These devices are expensive high-end test equipment, designed to accurately measure the absolute power of transmitters (e.g., for certification), or discover bugging devices that are transmitting in esoteric bands. Spectrum analyzers only measure the power of transmissions in the

frequency domain, they do not collect time-domain signals. Therefore, they cannot be used to perform signal analysis such as signal classification. For example, signals operating on the same frequency cannot be differentiated (e.g., in 2.4 GHz, antiquated DSSS 802.11b looks the same as modern OFDM 802.11g/n) on spectrum analyzer displays.

FMCW-BASED SPECTRUM SENSORS: As an improvement over spectrum analyzers which can only observe power, recent work by Cheema et al. [9] introduced receivers that can rapidly sweep over the spectrum to capture short time-domain samples across the spectrum. Their work is a proof-of-concept that demonstrates, with ideal hardware—namely, a costly signal generator—it is possible to perform high time resolution spectrum occupancy detection. This work inspired us to look into a practical modification for off-the-shelf radios that can make them sweep. However, unlike SweepSense, Cheema et al. only demonstrate using these samples to improve the time resolution of the spectrum occupancy. SweepSense is the first to demonstrate how to unsweep the samples to successfully perform signal analysis across GHz of spectrum, only with short captures of each band (Section 5). Prior to SweepSense, wide-bandwidth signal analysis was only considered possible with wide-bandwidth radios.

In summary, SweepSense demonstrates that narrow-bandwidth radios can be modified—with only the addition of an analog ramp generator fed to the VCO's tuning input—to create a rapidly sweeping radio. SweepSense also introduces a novel algorithm to unsweep distorted samples captured by modified off-the-shelf radios (Section 4). SweepSense also demonstrates that these unswept samples can still be used to perform rigorous signal analysis such as signal classification (Section 5).

INTELLIGENT SCANNING FOR SDR-BASED SENSORS: SpecInsight [37] improves the time resolution of spectrum sensing with narrow-bandwidth SDR's ($\sim$25 Msps) by intelligently scheduling when bands should be tuned into. Those that contain continuous transmitters (e.g., FM Radio) or predictable transmitters (e.g., airport RADAR) are tuned into infrequently, thereby improving the time resolution of narrow-bandwidth spectrum sensors. SpecInsight is complementary to SweepSense because it can use their band selection algorithm to intelligently select when to sweep particular bands. Therefore, other intelligent scanning algorithms [43, 26, 44, 25] can also be integrated into SweepSense to improve its time resolution.

SUB-NYQUIST SPECTRUM CAPTURE: Similar to SweepSense's goal of modifying off-the-shelf radios to operate across a wide bandwidth, prior work [4, 18] demonstrates that an off-the-shelf SDR can sample outside of their Nyquist bandwidth by removing the anti-aliasing filter on the RF frontend. However, these techniques assume that spectrum is sparsely occupied, and make use of specialized techniques like sparseFFT [16, 13, 17], or compressed sensing [4, 2, 10, 39]. SweepSense does not make such assumptions about the power and frequency of the transmissions in the spectrum, However, given that these systems are built on the same inexpensive SDRs as SweepSense, we might be able to increase our instantaneous bandwidth by sampling at sub-Nyquist rate while sweeping.

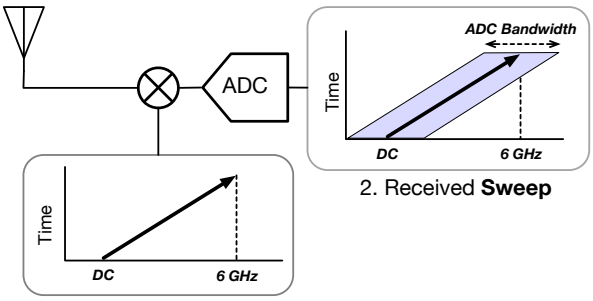

Figure 2: By chirping a receiver's Local Oscillator, it will rapidly sweep the spectrum.

3 Making Off-the-Shelf Radios Sweep

In this section, we describe how we modify the oscillator in off-the-shelf radios so they can rapidly sweep across several GHz. Fig. 2 shows an overview of the operation of a SweepSense receiver. To make the radio sweep, we modify the behavior of the radio's Local Oscillator (LO)— the device that tunes a radio into a particular frequency—to rapidly increase its frequency (chirp).

First, we describe how the LO in a radio can be modified to make it chirp. Then, we explain how to perform this modification on a USRP N210 SDR—a common off-the-shelf SDR with a wide tuning range.

3.1 How to make an LO chirp

To understand how to modify the LO to chirp, we must first explain how the LO operates in a radio. The LO is the hardware component in a radio that generates a tone which gives the receiver the ability to tune into a specific frequency. The tone from the LO is mixed with the amplified signal from the antenna to change the frequency of the received radio frequency (RF) signal and downconvert it to baseband. The baseband signal is then filtered and sampled by an ADC, and the raw digital samples are transferred to the host. Radios with a wide tuning range (e.g., SDRs) are built with a special LO that can generate tones across a wide frequency range; these LOs are called "wideband frequency synthesizers". For instance, the MAX2870 [29] frequency synthesizer on

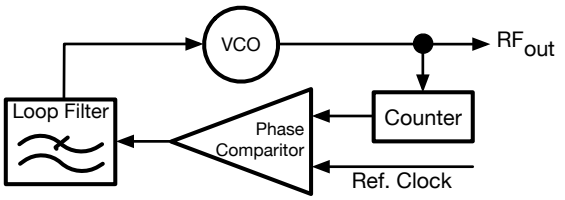

(a) **Normal config:** Closed-loop tuning

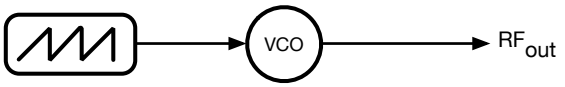

(b) **Sweeping config:** Open-loop chirping

Figure 3: Replacing a PLL-based LO's tuning feedback loop with a sawtooth waveform makes it sweep.

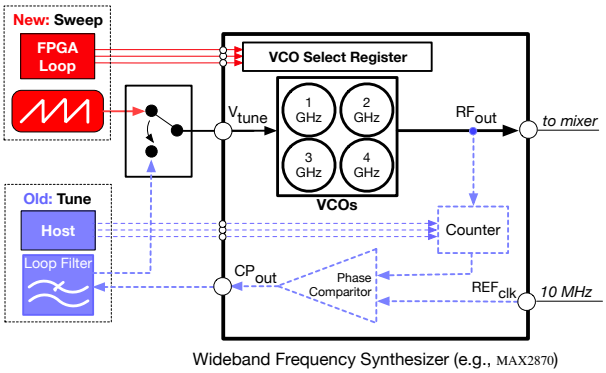

Figure 4: Modifications required to implement SweepSense on a COTS wideband frequency synthesizer.

the USRP CBX daughter card can generate tones ranging from 23.5 MHz to 6 GHz.

A wideband frequency synthesizer is commonly implemented using a Voltage Controlled Oscillator (VCO) in a highly integrated Phase Locked Loop (PLL). A simplified block diagram of a PLL is shown in Fig. 3 (a). The input voltage of the VCO determines its output frequency, and the PLL serves as the feedback loop that maintains control over the VCO input voltage to generate a fixed frequency tone. The feedback loop is driven by a phase comparator that compares the phase of the VCO output (divided by the counter), and the reference clock. The difference in phase is an indirect measure of the frequency error between the desired VCO output and its actual value. The external passive low-pass "loop filter" then filters the phase comparator output. The loop filter output drives the VCO input voltage, completing the control loop and "lock"ing the VCO output to the desired frequency. The loop filter characteristics and cutoff frequency determine the stability and accuracy of the frequency lock. Each time the frequency synthesizer is requested to generate a different frequency output, the PLL takes 10–100 μs to lock, during which the radio is temporarily offline. It is this repeated downtime that SweepSense avoids by making the PLL sweep continuously across a wide frequency range.

There are two parts of such an LO design that make them amenable to sweeping (1) the ability to control output frequency by adjusting the input voltage to the VCO, and (2) the customizable loop filter that is implemented with external passive components.

An LO can be modified to sweep by first disconnecting (by desoldering) the loop filter components, giving direct access to the VCO control input. Then, the now-unconnected VCO control input is connected to an externally generated sawtooth voltage. As the sawtooth signal repeatedly ramps its voltage, the VCO to repeatedly ramps its output frequency. As a result, the VCO output is a series of chirps (the modification is shown in Fig. 3 (b)). However, removing the feedback loop introduces several new challenges that we describe and address in Section 4.

Additionally, the wideband frequency synthesizers in off-the-shelf radios are particularly amenable to sweeping for spectrum sensing because they contain a bank of VCOs[2], each of which has a smaller frequency range (e.g., 100 MHz) that, put together, contribute to the LO's wide frequency range (depicted in Fig. 4). This modular construction makes such synthesizers much less expensive as compared to a single VCO synthesizer that has comparable tuning range. Also, being able to select which VCOs are used is important for frequency planning, such as skipping entire VCO bands that do not have active transmitters (SpecInsight [37]). Many modern frequency synthesizers (like the MAX2870) provide an explicit control register to select a particular VCO. For such synthesizers, SweepSense can implement fine-grained VCO selection and sweep with virtually no delay introduced due to the selection process.

3.2 Proof of Concept: Sweeping USRP

We now describe the complete modification that makes the commonly available USRP N210 SDRs sweep. We demonstrate that these modifications are general by performing them on three popular RF frontends for the USRP: the WBX and SBX that have an older Analog Devices synthesizer, and the CBX that has a modern Maxim synthesizer. We also believe it is compatible with the HackRF One that has a modern synthesizer from Qorvo. There are two aspects to this modification: (1) a hardware modification to disconnect the VCO feedback loop and replace it with a sawtooth signal and (2) an FPGA

[2]VCOs are implemented as a set of LC circuits (VCO cores) each of which can switch in a set of varactors (bands) depending on the desired frequency range

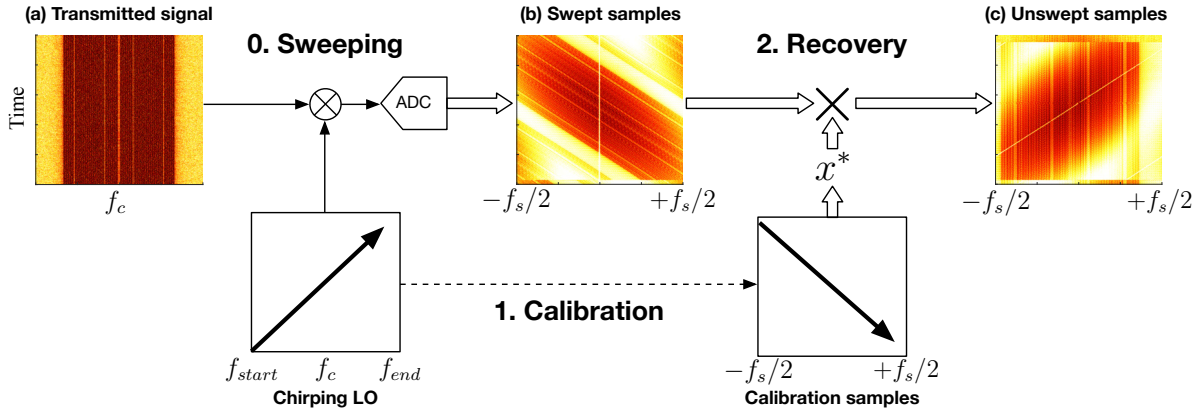

Figure 5: Illustration of the the signal captured by SweepSense at different stages in the receiver.

logic modification that makes the PLL cycle through its VCO bands, and generates the sawtooth waveform to sweep each VCO. Fig. 4 provides a visual overview of both modifications. The hardware schematics and Verilog needed to implement SweepSense will be made open source at the time of publication.

The hardware modification is straightforward for someone with surface mount soldering experience: you need to remove a single passive component from the SDR's receiver RF frontend, and in its place connect a wire that connects to one of the USRP's auxiliary Digital to Analog converters.

The FPGA's frequency synthesizer control logic modification removes all tuning logic, and in its place we add logic to iteratively loop through the selected VCO bands. Also, a new logic module is added to generate a sawtooth waveform and send it to the auxiliary DAC. These two logic modules are designed to operate in sync with the USRP's ADC sampling clock. This is required because unsweeping the samples requires knowing the configuration of the PLL, including its approximate tuning voltage, while the USRP is collecting each sample.

4 Unsweeping the Samples

Unlike a standard radio which samples with a local oscillator tuned to a fixed center frequency, SweepSense samples are distorted because they are captured while the center frequency is rapidly increasing. To aid in understanding the effect of a chirping local oscillator on captured samples, we begin with a primer on downconversion.

For a received signal $x(t)$ centered at frequency f_c as shown in Fig. 5(a), a standard fixed frequency direct IQ downconversion can be modeled as:

$$x_b(t) = x(t) \times e^{-j2\pi f_c t} \tag{1}$$

Where $x_b(t)$ is the downconverted signal (before baseband filtering) and f_c is also the frequency of the oscillator. In SweepSense, the oscillator frequency varies with time as $f(t)$. In our implementation, $f(t)$ monotonically increases with time (chirp). Therefore, similar to Eq. 1, a chirp direct IQ downconversion can be modelled as:

$$x_c(t) = x(t) \times e^{-j2\pi f(t) t} \tag{2}$$

This equation shows how sweeping introduces a significant change to the received signal: the frequency with which $x(t)$ is multiplied in SweepSense changes at every instant, and is offset from a fixed frequency oscillator at f_c by $f_c - f(t)$. Since $f(t)$ monotonically increases with the sawtooth waveform connected to the VCO tuning input, the frequency offset continuously decreases with time as shown in Fig. 5(b). The problem is, standard digital signal processing techniques rely on the assumption that the signal is fixed around a constant frequency at all times; therefore, these techniques can not be applied directly to the swept samples captured by SweepSense.

Undoing the sweeping effect requires removing the time-varying frequency offset $f_c - f(t)$ from SweepSense samples at time t, for which $f(t)$ is required. We call this process of undoing the sweeping effects "unsweeping". Unsweeping involves two steps:

1. Calibration: First, we extract the effect of sweeping ($f(t)$) by sending a known signal: we measure the frequency offset $f_c - f(t)$ introduced by SweepSense at time t.

2. Recovery: Then, we reverse the effect of sweeping by removing the offset $f_c - f(t)$ from the samples captured with SweepSense.

In summary, this method measures the sweeping center frequency, and uses it to recover signals as if they were captured at a fixed frequency.

4.1 Calibration

Why is calibration difficult?

The VCO's frequency increases as the voltage of the sawtooth waveform increases. Intuitively, one may expect that the VCO's frequency is directly related to the input voltage. However, this is not true for an open loop VCO (Section 3). An open loop VCO's frequency does not have a linear relationship with the input voltage : it is also dependent on temperature and other environmental conditions. However, we do know that the frequency increases monotonically as the input voltage increases. Therefore, to calibrate the VCO, we need to find another way to measure the center frequency $f(t)$ of SweepSense at each time instant in a sweep.

Insight and solution

Our insight is, we can calibrate the VCO by sweeping while capturing a tone transmitted at a known frequency. We measure the value of $f(t)$ by sending a tone at frequency f_c ($x_f(t) = e^{j2\pi f_c t}$) and collecting the received samples $x_{cal}(t)$ after the chirped direct downconversion, i.e., $x_{cal}(t) = e^{j2\pi(f_c - f(t))t}$ (Equation 2) as shown in Fig. 6. In summary, we directly capture the varying oscillator frequency in $x_{cal}(t)$. The implementation details of our calibration process appear in Section 6.

Calibration needs to be repeated at many reference tones due to the effect of the narrow-band radio's low-pass baseband filter on $x_{cal}(t)$. This filter suppresses the parts of $x_{cal}(t)$ whose frequencies lie outside the interval $[-F_s/2, F_s/2]$. Since the instantaneous frequency of $x_{cal}(t)$ is $f_c - f(t)$, it is detectable at time t only if $|f(t) - f_c| \leq F_s/2$. Therefore, for a specific tone, we can only calibrate the VCO behavior between $[f_c - F_s/2, f_c + F_s/2]$ using a tone of frequency f_c. To calibrate VCO's behavior at an arbitrary frequency interval $[f_{start}, f_{end}]$, we divide the calibration into chunks of bandwidth F_s and transmit a different reference tone for each chunk. Consecutive tones are each separated in frequency by F_s starting from $f_{start} + F_s/2$. We collect the received samples for all $x_{cal}(t) = e^{j2\pi(f_c - f(t))t}$ where $f_c = f_{start} + k * F_s$ where k = 1,2,..$(f_{end} - f_{start})/F_s$. This produces calibration data for the behavior of the VCO across the entire sensing bandwidth. This process only needs to be redone when temperature and environmental conditions change significantly.

4.2 Recovery

Next we describe how to use the data gathered in the calibration process to remove the time-varying frequency offset ($f_c - f(t)$). Recall that the downconversion in SweepSense VCO can be modeled as multiplying a chirp with the received signal. We observe that the frequency

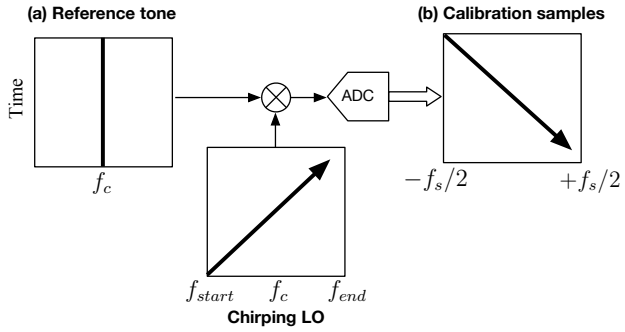

Figure 6: VCO behavior over F_s bandwidth is calibrated by sweeping over a reference frequency and collecting the samples.

of the calibration samples also varies similarly with time, motivating a similar multiplication to remove the effect of the chirp. Indeed, by multiplying the swept samples with the complex conjugate(.*) of the calibration samples $x_{cal}(t)$, it cancels out the frequency offset. Mathematically, the effect of sweeping cancels as follows:

$$\begin{aligned} x_c(t) \times x^*_{cal}(t) &= [x(t) \times \cancel{e^{-j2\pi f(t)t}}] \times \cancel{e^{j2\pi f(t)t}} e^{-j2\pi f_c t} \\ &= x(t) \times e^{-j2\pi f_c t} \end{aligned} \quad (3)$$

This process converts a chirped direct downconversion to the corresponding fixed frequency downconversion as in Equation 1. Therefore, signals are recovered as if they were received by a standard fixed frequency receiver. We evaluate the performance of unsweeping in Section 7.

Fig. 7 shows an example of signals captured between 2.380 GHz and 2.480 GHz after their recovery using the calibration data. In this capture, we observe multiple OFDM packets centered at 2.412 GHz (even an acknowledgment packet around 400 μsec) and a Bluetooth packet at 2.428 GHz. Unlike FMCW spectrum sensors which can only detect signal energy, SweepSense can capture short intervals of the time-domain samples of the transmitted signal. These samples enable SweepSense to distinguish different transmissions, even when they have a similar center frequency and bandwidth. Unsweeping therefore is an improvement to prior high-speed sweeping spectrum sensing architectures (Section 2).

5 Analysis and Inference

In this section, we describe a method to detect modulation scheme and protocol type from swept samples. Conventional detection algorithms for signal classification fixed frequency spectrum sensors rely on capturing a significant portion of the transmission, sometimes even requiring protocol-specific preambles [27]. However, SweepSense only captures a small number of samples for each frequency band. Hence, it is unlikely that

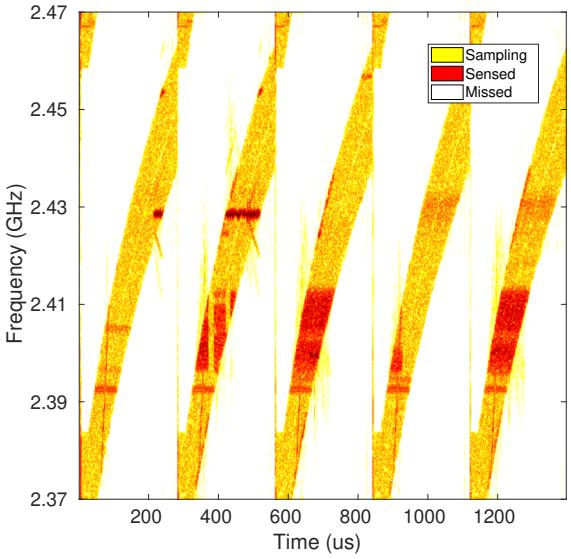

Figure 7: Example of ISM-band transmissions recovered from swept samples.

it will capture a preamble. Also, the open-loop operation of VCO during sweeping introduces additional noise into the signal, making it difficult to perform signal classification.Therefore, we designed a classification algorithm that is resilient to noise, and works even with only a short capture of the signals.

Our algorithm is inspired by cyclostationary analysis [12]. The basic premise behind cyclostationary analysis is that every human-made signal has inherent periodicity associated with it. This periodicity is unique to every protocol, independent of implementation or hardware used. It also can serve as a fingerprint for inference [19]. For example, in WiFi-OFDM, the cyclic prefix (CP) repeats at the start and end of a symbol. SweepSense's key insight is that this periodicity is retained even when we receive a small portion of transmission filtered in time and frequency. Cyclostationary functions evaluate this periodicity as correlations in time and frequency domains. SweepSense uses these cyclostationarity signatures to build reliable ML models for signal classification. For all our analyses, we use two second-order cyclostationary functions: the Cyclic Autocorrelation Function (CAF) and the Spectral Correlation Function (SCF).

If $x[n]$ is the received signal, the CAF estimator is calculated as follows [8]:

$$R_x^{\alpha}(\tau) = \sum_{n=-\infty}^{\infty} x[n]\,[x^*[n-\tau]]\,e^{-j2\pi\alpha n} \qquad (4)$$

The CAF is maximized when the choice of delay (τ) is equal to the time between consecutive repeating patterns in $x[n]$. This causes them to align in the correlation. These maxima occur periodically along n, and the

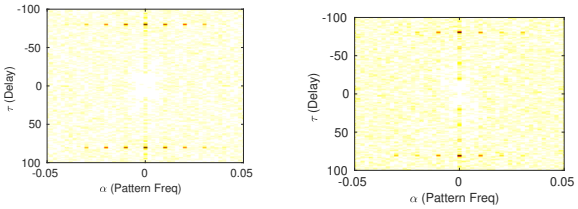

(a) Fixed freq. receiver (b) SweepSense

Figure 8: The CAF is visible in SweepSense captures.

term $e^{-j2\pi\alpha n}$ is a transform that brings out the frequency (α) of this periodicity. α may be interpreted as the frequency of repetition of hidden patterns, defined as the pattern frequency. Therefore, CAF peaks at values of τ and α that correspond respectively to the time period and repetition frequency of patterns in $x[n]$. The CAF is particularly useful in analyzing signals like OFDM with repetitive patterns in time (i.e., cyclic-prefixes [38]). The SCF is the Fourier transform of the CAF over τ, making them equivalent representations due to the unitary nature of the transform. The SCF peaks for the same values of α as the CAF and frequency f is the fourier dual of delay τ. The SCF can be efficiently computed due to its representation using FFTs as described below.

Consider L consecutive discrete time windows of $x[n]$, each of length N samples. $X_{lN}(f)$ is the FFT of $x[n]$ for the l^{th} time window. The time-smoothed SCF estimator for this signal is calculated as follows [8]:

$$S_x^{\alpha}(f) = \frac{1}{LN}\sum_{l=0}^{L-1} X_{lN}(f)X_{lN}^*(f-\alpha) \qquad (5)$$

As an illustration, Fig. 8(a) shows the CAF plot of WiFi-OFDM. The x-axis represents pattern frequency (α) and the y-axis represents delay (τ). WiFi symbols are 80 samples long (of which 16 are CP) at 20 MHz sampling rate. Since we sample at 25 Msps, we get 100 samples per symbol (of which 20 are CP). Notice that the CAF peaks at a τ =80 samples and α =0.01 (normalized to 25 Msps). We also observe peaks in the SCF plot (not shown) at the same α values. Patterns such as these occur in every protocol, we do not need to capture the entire packet to identify them. Indeed, we see in Fig. 8(b) that the CAF of the unswept samples of WiFi-OFDM also exhibits the peaks at same points as the fixed frequency capture. The CAF and SCF are robust due to their highly signal selective nature, magnifying the signal's natural patterns while averaging and suppressing distortions introduced due to sweeping.

For ML-based classification, we extract CAF and SCF features from the unswept signal at a set of precomputed values of α, τ and f. Specifically, we to include values that are at the expected peaks for the protocols that we seek to detect.

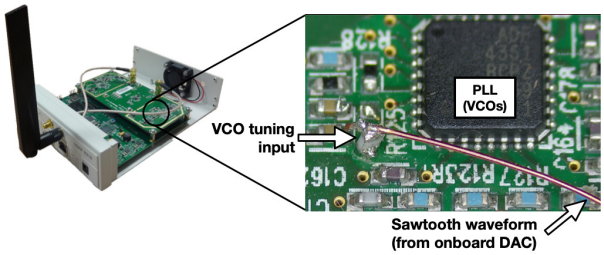

Figure 9: SweepSense requires a single-wire modification to the USRP's RF frontend to make the PLL chirp.

6 Implementation

Our hardware setup for SweepSense uses a standard off-the-shelf USRP N210 SDR. We implement the LO modification as discussed in Section 3 on both the CBX daughter card which supports 1.2 GHz to 6 GHz and SBX daughter card which supports 400 MHz to 4.4 GHz (shown in Figure 6).

We then make the following modifications to the FPGA logic on the USRP. The voltage ramp used to control the VCO is generated using the AUX-DAC on the daughtercard, which is controlled by the FPGA. Special care is taken to ensure that the voltage generated by the AUX-DAC on the USRP is time synchronised with the baseband ADC samples. The added logic also selects the PLL's VCO band and RF divider. The user can configure the sweeping bandwidth (VCO band and RF divider selection) and sweep rate (sawtooth voltage ramp slope) from the GNURadio python environment.

For our observations in the ISM band, we use a standard 2.4 GHz omnidirectional antenna, and for our wideband captures, we use a discone antenna mounted on the roof of the CSE building at UC San Diego. We operate the USRP at a sampling rate of 25 MSps with 16-bit resolution. For the evaluation, the captured samples are streamed, stored on the PC and processed offline.

Calibration and Recovery

The calibration process is as follows: SweepSense transmits a reference tone from the (unmodified) transmit chain of USRP. It receives the tone with the (modified) sweeping receive chain indirectly from leakage between the transmitter and recevier RF paths[3]. To calibrate across the entire sensing bandwidth, we repeat this process with tones separated by the sampling bandwidth (Section 4.1). For example, we need to run the calibration process 200 times when the sampling bandwidth is 25 MHz and the sweeping bandwidth is 5 GHz. In each of these files consisting calibration data for a different tone, the samples where the sweeping of a VCO band starts and ends is deterministic since the voltage input to VCO is synchronized with start of ADC sampling. Further, since these tones are separated by F_s the time intervals during which they are received are non-overlapping. Therefore we can combine the calibration data from these multiple tones by just adding the data from each file.

Periodic re-calibration may be necessary due to frequency drift of the VCO, particularly when the ambient temperature significantly changes (details in Section 7.2). However, re-calibration only requires performing one sweep over each of the reference tones. For example, calibrating at a sweeping bandwidth of 5 GHz and rate of 125 μsec/100 MHz only requires 6.25 milliseconds of downtime.

SweepSense recovers the time-domain samples from the swept samples in real time. This is feasible because recovery only requires performing conjugate multiplication of the swept samples with the calibration samples (Section 4.2).

[3]This is inspired by the USRP's use of TX/RX leakage to calibrate for I/Q imbalance.

7 Evaluation

To evaluate the performance of SweepSense as a spectrum sensor we first evaluate the performance of SweepSense with several high-time-resolution spectrum sensing case-studies that normally would require a wide-bandwidth spectrum sensor. Then, we evaluate the limitations of SweepSense with several micro-benchmarks.

We selected the case studies based on the results of a sample full spectrum (0–6 GHz) capture that we performed in the lab. Although there were many occupied bands in this capture, we observed that the 2–3 GHz spectrum was the most dynamic (shown in Figure 10) due to nearby WiFi, Bluetooth, and LTE deployments. In the ISM band (2.4 GHz), we demonstrate that we can detect and classify diverse protocols. In addition, we show how SweepSense can monitor the load on multiple LTE base stations (1.9–2.2 GHz) simultaneously. We conclude the case studies by evaluating the performance of SweepSense as an Environment Sensing Capability (ESC) sensor for the newly shared 3.5 GHz Citizens Broadband Radio Service (CBRS) spectrum [40].

The micro-benchmarks evaluate the frequency distortion and signal to noise ratio (SNR) loss due to the sweep and unsweep processes, and a demonstration of frequency stability across sweeps.

In summary, our evaluation contains the following the results:

- Protocols can be classified based on unswept samples containing partial packets or a few symbols, usually requiring only 25 μs to classify the signal types in contrast to typical full packet lengths 1–10 ms, an improvement of over 40 $\times$.

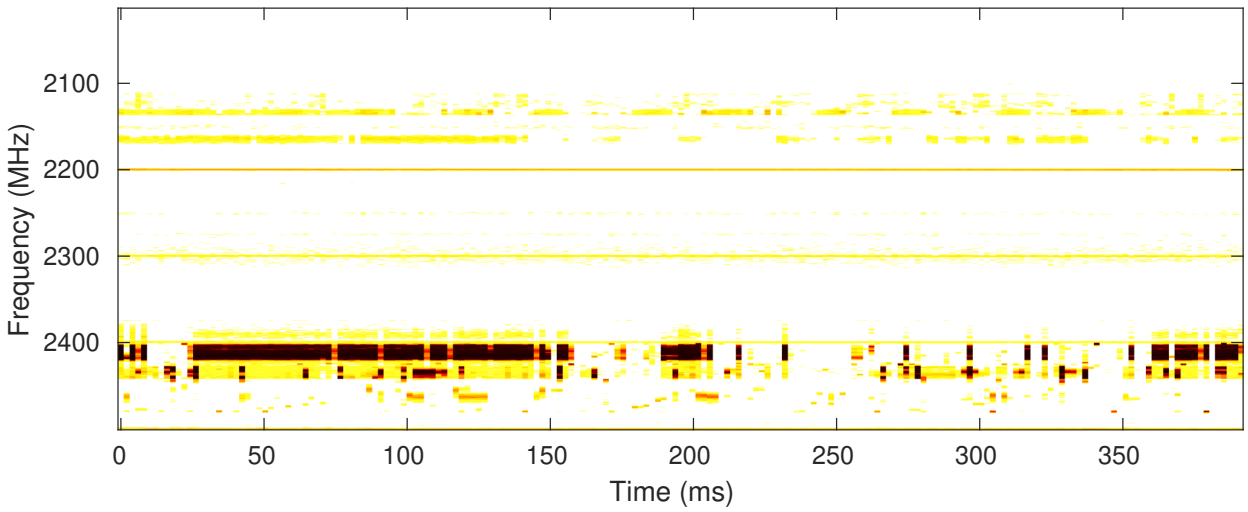

Figure 10: Example of transmissions between 2 and 2.5 GHz captured by SweepSense every 2 ms.

- In contrast to a standard CBX receiver, which takes 500 μs to monitor 100 MHz (125μs to capture and retune four times), SweepSense can do it in 125 μs, a 4$\times$ improvement.
- Useful information such as channel utilization can be extracted with 1 ms resolution in highly dynamic and disjoint LTE distributed in a band of 200 MHz.
- Incumbent sensing can be reliably performed over 200 MHz of bandwidth for use in spectrum sharing architectures like CBRS.
- The loss of quality in received samples due to a free-running VCO and the unsweeping mechanism can be characterized and do not limit the use of SweepSense as a spectrum sensor

Our evaluation hardware setup is as described in the previous section. We select VCO bands and sweep rates that best suit the evaluation requirements. In situations where comparisons are required, we use an oracle to provide the ground truth The oracle is an unmodified USRP (CBX frontend) synchronized with SweepSense using a MIMO cable. The oracle USRP is tuned to a particular frequency, while the SweepSense USRP continuously sweeps multiple bands. We then repeat the experiments while cycling the oracle through all of the relevant frequency bands.

7.1 Case Studies

7.1.1 ISM Protocol Classification

In the first case study we evaluate the performance of SweepSense in differentiating between four common protocols in the ISM band: WiFi-OFDM (802.11g/n), WiFi-DSSS (802.11b), Bluetooth (BLE), Zigbee (ZB), and no transmission (Gaussian noise). These protocols are diverse in their bandwidth, modulation scheme, and behavior. Both WiFi-DSSS and WiFi-OFDM are relatively wideband but have the same bandwidth (20 MHz) and channel allocation. [20] BLE and ZB are relatively narrowband (2 MHz), and have overlapping, but different channel allocation, making the classification process more difficult [21, 22].

We used a two-level classifier to distinguish between the various protocols. The first level differentiates between narrowband and wideband signals using the Power Spectral Density (PSD). The second level then implements an SVM classifier for the wideband signals and a single layer neural network for narrow band signals [6]. Both of these classifiers take as input vectors the SCD and CAF of the unswept samples within each sweep. For wideband signals, CAF vectors are obtained at cyclic frequency shifts of $k*0.01$; and for narrowband signals, they are obtained at cyclic frequency shift of $k*0.0025$. The classifiers were trained using ground truth captures of each protocol captured over the air. The ground truth signals were generated using relevant MATLAB toolboxes or standard compliant scripts, and included signals at a wide range of SNRs. The first classifier (CLASSIFIER_1), differentiates between transmission (noise), ZB, and BLE. The second classifier (CLASSIFIER_2) differentiates between no transmission (noise), WiFi-OFDM, and WiFi-DSSS.

Classification accuracy is used as the primary metric in this evaluation, and is calculated as: the number of sweeps that were classified correctly, divided by the total number of sweeps where the signal was present. We performed the evaluation with a SweepSense receiver capturing signals over-the-air that we transmitted across the

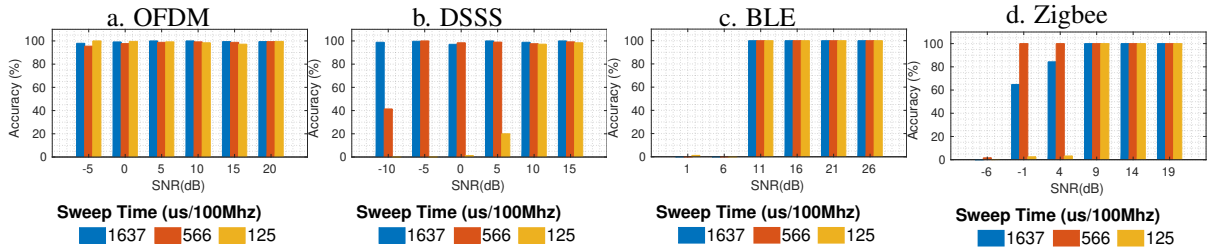

Figure 11: Classification accuracy for ISM protocols across SNRs.

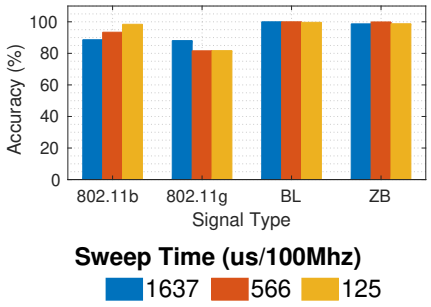

Figure 12: Classification accuracy for ISM protocols across multiple test locations.

entire 100 MHz wide 2.4 GHz ISM band. In each experiment, we transmit the ground truth signals containing a mix of protocols, and they are received simultaneously by co-located SweepSense and oracle USRP receivers. The classifier then operates on the unswept samples from the SweepSense receiver, and the ground truth samples from the oracle USRP. We then calculate the classification accuracy. The receiver setup is then moved around the lab to capture data from multiple locations.

Figure 12 shows the classification accuracy across all four protocols. The average classification accuracy for signals with the highest transmit SNR, across all protocols, is 95%. While operating on the fastest sweep rate of 125 μs per 100 MHz, we repeat the experiment while varying the sweep rate and transmit SNR to understand the classification accuracy's dependence on these parameters. Figure 11 shows that SweepSense classification accuracy is high for signals with decodable SNR even at the fastest sweep rate of 125 μs per 100 MHz: CLASSIFIER_1 can detect and classify signals with 95% accuracy at even low SNR with sweep as fast as 125 μs per 100 MHz. CLASSIFIER_2 can detect and classify signals with 90% accuracy at even low SNR with a sweep as fast as 125 μs per 100 MHz.

We note that the noise suppression properties of cyclostationary analysis enables us to correctly classify signals even when they're sometimes below the noise floor. The accuracy drops as the rate of sweep increases. We see that the drop in accuracy is because faster sweep rates lead to a smaller number of samples (the fastest sweep yields only 3125 samples in every 100 MHz). It also leads to larger distortions, both of which negatively affect cyclo-stationary signatures. It should be noted that these signatures are preserved at lower sweep rates, despite the frequency distortions.

7.1.2 LTE Channel Utilization

The LTE bands are allocated to specific service providers, but even within a service provider, the bands are across a wide frequency range in the spectrum. Also, LTE base stations schedule traffic at a millisecond granularity. Therefore, monitoring the load across many LTE base stations demonstrates SweepSense's ability to capture time dynamics of signals across a wide bandwidth. Specifically, we show that SweepSense can simultaneously monitor the load of a set of disjoint LTE downlink channels (with a total bandwidth of 75 MHz), spread over the 1.9 GHz and the 2.1 GHz bands.

Our experimental setup is as follows: we connect the SweepSense receiver to a wideband discone antenna on the roof of the building. SweepSense is configured to sense 1.9 GHz to 2.1 GHz spectrum in three sweeps, each is 80MHz at the rate of 375 μs per 100MHz. We captured several seconds of sweeps during a peak hour in the evening.

Since the LTE protocol only puts energy on subcarriers when downlink traffic is transmitted, the energy of each subcarrier directly correlates with the downlink channel usage [1]. Therefore, we use a short-term Fourier transform on the unswept samples and report load as average power levels detected in the respective bands. The maximum power level obtained over all our experiments is used as the normalization factor to obtain the power corresponding to the maximum load. Fig. 13 shows a snapshot of simultaneously measured load of five LTE base stations with 0.9 ms granularity (less than the scheduling interval) per LTE base station. Surprisingly, even at peak hours, the load across base stations is very uneven.

7.1.3 CBRS ESC Sensor

The FCC requires spectrum sensing in the CBRS band to detect and avoid interfering with incumbent radar transmissions. Highly reliable ESC sensors that moni-

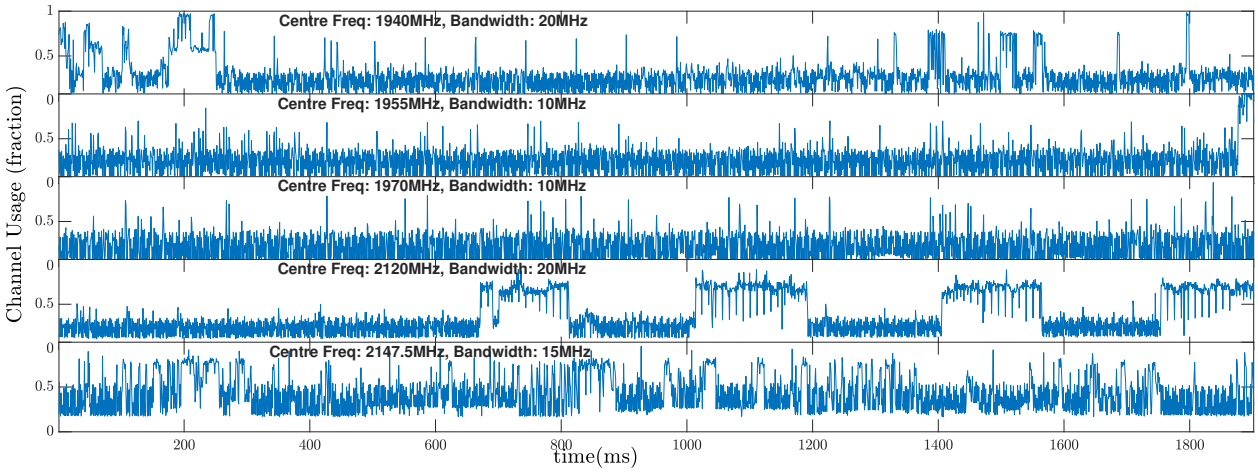

Figure 13: Downlink channel occupancy of five LTE base stations as observed simultaneously by SweepSense.

tor the entire spectrum for incumbent transmissions, are arguably one of the most critical parts of the rules for using the CBRS spectrum [40].

We evaluate the capability of SweepSense as an ESC sensor. Our experiment is to detect the "Bin 1 Lite" radar waveform as per the official testing and certification procedures for ESC sensors [11]: this radar type closely resembles widely deployed maritime pulse radar. We use MATLAB to generate the radar signals and add Gaussian noise (GN) according to the specified relative power levels in [11]. The samples are transmitted to the SweepSense USRP with a Vector Signal Generator (Keysight N5182B) at calibrated power levels. The signal generator is directly connected to the SweepSense receiver with RF coax. SweepSense is configured to sweep 3480 MHz - 3680 MHz every 1.3 ms. We sweep the spectrum multiple times within one radar burst interval, increasing chances of detection. In each experiment, we initiate the SweepSense capture for 10 seconds and then trigger the signal generator ten times. Our sensing algorithm declares radar events based on peaks in the short term Fourier transform of the unswept signal. Since the SweepSense USRP is not designed to have a low noise floor, the actual power levels used in this study are 9 dB/MHz higher (-80dBm/Hz for radar pulses and -100dBm/Hz for GN) than the respective values in [11].

Table 1 summarizes the radar detection performance of SweepSense. We observe that SweepSense can achieve 99.5% accuracy with a very simple receiver. Added to this, we also demonstrate that SweepSense can function as an ESC sensor over *double* the required bandwidth, motivating broader spectrum sharing applications in the future. In summary, SweepSense is effective in detecting fleeting signals (e.g., radar).

Radar Type	Pulse Width (μs)	Pulses per second	Pulses per burst	Detection accuracy
Bin 1 Lite	0.8	1000	19	99.5% (398/400)

Table 1: ESC radar classification accuracy

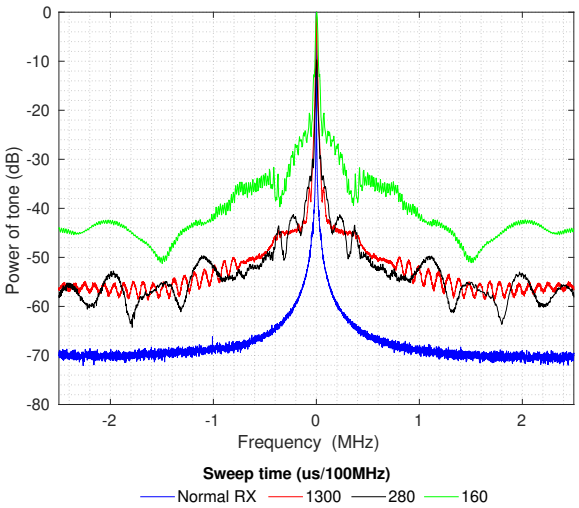

Figure 14: PSD characteristics of a fixed tone captured using SweepSense across multiple sweep rates.

7.2 Micro benchmarks

Frequency stability and phase noise are typical measurements used to characterize radios. Since the sweep–unsweep process recovers fixed frequency samples, we can benchmark the performance of SweepSense using

these standard metrics. We isolate the loss in performance due to sweeping by comparing with the performance of un-modified USRP SDR radios. In each of these evaluations, we connected a signal generator that outputs a single frequency tone into both the SweepSense USRP and the oracle USRP through identical RF paths. SweepSense sweeps the relevant band and uses pre-captured calibration data to obtain the unswept samples at different sweep rates.

Since the VCO in SweepSense is operating in open-loop mode, we observe a frequency drift over time (shown in Fig. 15). The rapid rise and subsequent settling of the frequency is due to the oscillator warming up and settling on its stable operating temperature after power-on. We observe that the settling time is consistent: it takes the same amount of time every time we power on the USRP ($\sim$1500 s), and it is also consistent across multiple VCO bands and sweep rates. Although the VCO takes many minutes to settle, this is only a one-time event at power-on and does not affect the performance of a SweepSense sensor after it has warmed up or switched bands.

Next, we characterize the performance of the SweepSense un-sweeping and noise distortion added due to un-sweeping compared to fixed frequency receiver. On a standard fixed frequency radio, the PLL reduces the phase noise of the VCO while it locks the frequency to the desired value. Since we removed the PLL lock loop for implementing SweepSense, it is essential to characterize the distortion created by the open loop VCO being controlled by the the sawtooth signal from an external DAC. All measurements are taken after the frequency drift settles. We compare Power Spectral Density of the unswept tone at different sweep rates against samples received by the oracle USRP in Fig. 14. We see that the phase noise floor rises by $\sim$10 dB for slower sweep rates and the skirt around 0 Hz starts increasing for higher sweep rates, compared to the oracle. An ideal response would have a clean tone with no skirt or spreading. Sweeping faster, therefore, comes at the cost of limited frequency resolution.

8 Limitation: SNR Loss and Inference

The phase noise of SweepSense will lead to a loss in signal quality. Phase noise is multiplicative noise, i.e., SNR loss due to phase noise depends on the signal strength of the transmission. If the transmission has 10 dB of SNR, i.e., the noise floor would be 10 dB lower than signal; then the effect of phase noise will be insignificant (less than 1 dB loss). Recall that the classification evaluation results demonstrate that even with a weak signal (e..g, 5 dB SNR), SweepSense can classify the 20 MHz OFDM signal with just a 25μsec capture (sampled at 25 Msps). This means that even with such high phase noise, the inference algorithms still perform well. In summary, SweepSense has high distortion due to phase noise, but even then it still performs well for signal detection.

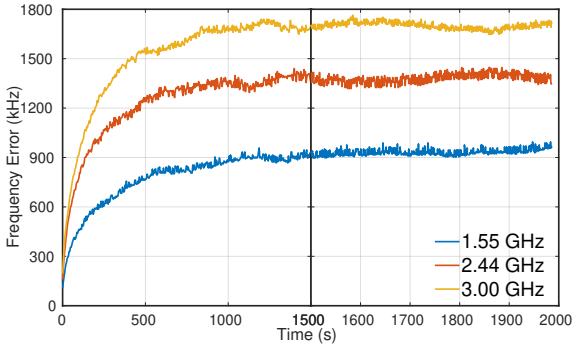

Figure 15: Frequency error in VCO output vs time of capture. The VCO reaches temperature stability in ($\sim$1500 s).

9 Conclusion

SweepSense presents the first spectrum sensor which can rapidly sweep the entire terrestrial spectrum with low-cost SDRs, while providing detailed measurements including transmitter classification and utilization. SweepSense achieves this by making a single-wire modification to the frontend of SDRs such as the USRP, allowing us to make this improvement to current deployments of USRP radios in multiple wide-scale deployments such as CityScape [35], and the Microsoft Spectrum Observatory [31].

In addition to spectrum sharing, SweepSense can be used for data mining, since communication signals are generated when humans, machines, and objects change their state. In the future we envision the community adding other spectrum analysis techniques beyond classifying communication protocol, namely transmitter localization.

References

[1] 3GPP Consortium. 3GPP Specification series 36. http://www.3gpp.org/dynareport/36-series.htm.

[2] O. Abari, F. Lim, F. Chen, and V. Stojanović. Why analog-to-information converters suffer in high-bandwidth sparse signal applications. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 60(9):2273–2284, Sep. 2013.

[3] Anritsu. MS2840A Spectrum Analyzer. https://www.anritsu.com/en-IN/test-measurement/solutions/ms2840a-066/.

[4] M. R. Avendi, K. Haghighi, A. Panahi, and M. Viberg. A NLLS based sub-nyquist rate spectrum sensing for wideband cognitive radio. *CoRR*, abs/1408.4544, 2014.

[5] P. Bahl, R. Chandra, T. Moscibroda, R. Murty, and M. Welsh. White space networking with Wi-Fi like connectivity. In *Proc. ACM SIGCOMM*, 2009.

[6] C. M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.

[7] A. M. Cavalcante, E. Almeida, R. D. Vieira, S. Choudhury, E. Tuomaala, K. Doppler, F. Chaves, R. C. D. Paiva, and F. Abinader. Performance evaluation of lte and wi-fi coexistence in unlicensed bands. In *IEEE Vehicular Technology Conference (VTC Spring)*, June 2013.

[8] chad spooner. Cyclostationarity blog. https://cyclostationary.blog/.

[9] A. A. Cheema and S. Salous. Digital FMCW for ultrawideband spectrum sensing. *Radio Science*, 51(8):1413–1420, Aug 2016.

[10] M. F. Duarte and R. G. Baraniuk. Spectral compressive sensing. *Applied and Computational Harmonic Analysis*, 35(1):111 – 129, 2013.

[11] F. H. Sanders, J. E. Carroll, G. A. Sanders, R. L. Sole, J. S. Devereux, and E. F. Drocella. "Procedures for laboratory testing of environmental sensing capability sensor devices" National Telecommunications and Information Administration, Techincal Memorandum TM 18-527". https://www.its.bldrdoc.gov/publications/3184.aspx, Nov. 2017.

[12] W. A. Gardner. The spectral correlation theory of cyclostationary time-series. *Signal Processing*, 11(1), July 1986.

[13] B. Ghazi, H. Hassanieh, P. Indyk, D. Katabi, E. Price, and L. Shi. Sample-optimal average-case sparse fourier transform in two dimensions. *CoRR*, abs/1303.1209, 2013.

[14] A. P. Goodson. A multi-function, broad band, high dynamic range RF receiver. Technical report, OneRadio, 2017.

[15] T. Harrold, R. Cepeda, and M. Beach. Long-term measurements of spectrum occupancy characteristics. In *New Frontiers in Dynamic Spectrum Access Networks (DySPAN), 2011 IEEE Symposium on*, pages 83–89. IEEE, 2011.

[16] H. Hassanieh, P. Indyk, D. Katabi, and E. Price. Nearly optimal sparse fourier transform. *CoRR*, abs/1201.2501, 2012.

[17] H. Hassanieh, P. Indyk, D. Katabi, and E. Price. Simple and practical algorithm for sparse fourier transform. In *Proceedings of the Twenty-third Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '12, pages 1183–1194, Philadelphia, PA, USA, 2012. Society for Industrial and Applied Mathematics.

[18] H. Hassanieh, L. Shi, O. Abari, E. Hamed, and D. Katabi. GHz-Wide sensing and decoding using the sparse fourier transform. In *Proc. IEEE Conference on Computer Communications (INFOCOM)*, 2014.

[19] S. S. Hong and S. R. Katti. DOF: a local wireless information plane. In *Proceedings of the ACM SIGCOMM 2011 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Toronto, ON, Canada, August 15-19, 2011*, pages 230–241, 2011.

[20] IEEE. IEEE 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications". (2016 revision). IEEE-SA. 14 December 2016. http://ieeexplore.ieee.org/document/7786995/.

[21] IEEE Standard. "802.15.4k-2013 - IEEE Standard for Local and metropolitan area networks–Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs)–Amendment 5: Physical Layer Specifications for Low Energy, Critical Infrastructure Monitoring Networks.". https://ieeexplore.ieee.org/document/6581828/.

[22] IEEE Standard. "IEEE Std 802.15.1–2005 – IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements Part 15.1: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks (W Pans)". leeexplore.ieee.org.

[23] M. H. Islam, C. L. Koh, S. W. Oh, X. Qing, Y. Y. Lai, C. Wang, Y. Liang, B. E. Toh, F. Chin, G. L. Tan, and W. Toh. Spectrum survey in singapore: Occupancy measurements and analyses. In *2008 3rd International Conference on Cognitive Radio Oriented Wireless Networks and Communications (CrownCom 2008)*, pages 1–7, May 2008.

[24] A. P. Iyer, K. Chintalapudi, V. Navda, R. Ramjee, V. N. Padmanabhan, and C. R. Murthy. SpecNet: Spectrum sensing sans frontières. In *Proc. Symposium on Networked Systems Design and Implementation (NSDI)*, 2011.

[25] H. Kim and K. G. Shin. Efficient discovery of spectrum opportunities with mac-layer sensing in cognitive radio networks. *IEEE transactions on mobile computing*, 7(5):533–545, 2008.

[26] H. Kim and K. G. Shin. Fast discovery of spectrum opportunities in cognitive radio networks. In *New Frontiers in Dynamic Spectrum Access Networks, 2008. DySPAN 2008. 3rd IEEE Symposium on*, pages 1–12. IEEE, 2008.

[27] K. Lakshminarayanan, S. Sapra, S. Seshan, and P. Steenkiste. Rfdump: An architecture for monitoring the wireless ether. In *Proceedings of the 5th International Conference on Emerging Networking Experiments and Technologies*, CoNEXT '09, pages 253–264, New York, NY, USA, 2009. ACM.

[28] M. Lopez-Benitez, A. Umbert, and F. Casadevall. Evaluation of spectrum occupancy in spain for cognitive radio applications. In *Vehicular technology conference, 2009. VTC Spring 2009. IEEE 69th*, pages 1–5. IEEE, 2009.

[29] Maxim Integrated. 23.5MHz to 6000 MHz Fractional Integer-N synthesizer/VCO. http://hforsten.com/third-version-of-homemade-6-ghz-fmcw-radar.html.

[30] MetaGeek. Wi-Spy and Chanalyzer. https://www.metageek.com/products/wi-spy/.

[31] Microsoft. Spectrum Observatory. http://observatory.microsoftspectrum.com/.

[32] OSCOR. Blue Spectrum Analyzer. https://reiusa.net/rf-detection/oscor-blue-spectrum-analyzer/.

[33] K. Qaraqe, H. Celebi, M. Alouini, A. El-Saigh, L. Abuhantash, M. Al-Mulla, O. Al-Mulla, A. Jolo, and A. Ahmed. Measurement and analysis of wideband spectrum utilization in indoor and outdoor environments. In *International Conference on Communications Technologies (ICCT 2010)*. Citeseer, 2010.

[34] S. Rayanchu, A. Patro, and S. Banerjee. Airshark: Detecting non-WiFi RF devices using commodity WiFi hardware. In *Proc. ACM Internet Measurement Conference (IMC)*, 2011.

[35] S. Roy, K. Shin, A. Ashok, M. McHenry, G. Vigil, S. Kannam, and D. Aragon. Cityscape: A metro-area spectrum observatory. In *Proc. IEEE International Conference on Computer Communication and Networks (ICCCN)*, 2017.

[36] S. Salous, N. Nikandrou, and N. Bajj. Digital techniques for mobile radio chirp sounders. *IEE Proceedings-Communications*, 145(3):191–196, 1998.

[37] L. Shi, P. Bahl, and D. Katabi. Beyond sensing: Multi-GHz realtime spectrum analytics. In *Proc. Symposium on Networked Systems Design and Implementation (NSDI)*, 2015.

[38] P. D. Sutton, K. E. Nolan, and L. E. Doyle. Cyclostationary signatures in practical cognitive radio applications. *IEEE Journal on Selected Areas in Communications*, 26(1):13–24, Jan 2008.

[39] Z. Tian and G. B. Giannakis. Compressed sensing for wideband cognitive radios. Technical report, MICHIGAN TECHNOLOGICAL UNIV HOUGHTON, 2007.

[40] U.S. Government. CFR title 47 section 96.67 Environmental Sensing Capability.

[41] M. Wellens, J. Wu, and P. Mähönen. Evaluation of spectrum occupancy in indoor and outdoor scenario in the context of cognitive radio. In *CrownCom*, pages 420–427, 2007.

[42] J. Xue, Z. Feng, and P. Zhang. Spectrum occupancy measurements and analysis in beijing. *IERI Procedia*, 4:295–302, 2013.

[43] S. Yoon, L. E. Li, S. C. Liew, R. R. Choudhury, I. Rhee, and K. Tan. Quicksense: Fast and energy-efficient channel sensing for dynamic spectrum access networks. In *INFOCOM, 2013 Proceedings IEEE*, pages 2247–2255. IEEE, 2013.

[44] T. Yucek and H. Arslan. A survey of spectrum sensing algorithms for cognitive radio applications. *IEEE communications surveys & tutorials*, 11(1):116–130, 2009.