

Spectrum Monitoring and Source Separation in POWDER

Boston C Terry
boston.terry@utah.edu
University of Utah
Salt Lake City, Utah

Alex Orange
University of Utah, Salt Lake City
Utah

Neal Patwari
Washington University in St. Louis
One Brookings Drive, St. Louis
Missouri

Sneha Kasera
University of Utah, Salt Lake City
Utah

Jacobus Van Der Merwe
University of Utah, Salt Lake City
Utah

ABSTRACT

Current software-defined radio systems enable transmission at nearly arbitrary frequencies, presenting the possibility of harmful interference to existing communication services when broadcasting over-the-air. The Platform for Open Wireless Data-driven Experimental Research (POWDER) provides software radios whose output can be amplified and transmitted over-the-air. POWDER must include a spectrum monitoring system that can identify users who are transmitting outside allowed frequency bands to ensure wireless spectrum license holders do not experience harmful interference.

Power amplifiers in the transmit signal path can create emissions at center frequency harmonics and other spurious emissions. A spectrum monitoring system, coupled with signal paths after all amplifiers in the transmit chain, can detect these emissions. However, incident radio frequency energy combines with the output signal, which is no longer buffered by the amplifier. Incident and transmitted signals must be separated and isolated. The monitor can then analyze the isolated transmitted signal for out-of-band energy. This paper presents a system that can achieve isolation and identify users that broadcast out-of-band.

CCS CONCEPTS

• **Networks** → **Network experimentation; Mobile networks; Wireless access points, base stations and infrastructure;** • **General and reference** → *Experimentation.*

KEYWORDS

spectrum sensing, source separation, wireless testbed

ACM Reference Format:

Boston C Terry, Alex Orange, Neal Patwari, Sneha Kasera, and Jacobus Van Der Merwe. 2020. Spectrum Monitoring and Source Separation in POWDER. In *14th International Workshop on Wireless Network Testbeds, Experimental evaluation & Characterization (WiNTECH'20), September 21, 2020, London, United Kingdom*. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3411276.3412192>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WiNTECH'20, September 21, 2020, London, United Kingdom

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8082-9/20/09...\$15.00

<https://doi.org/10.1145/3411276.3412192>

1 INTRODUCTION

The Platform for Open Wireless Data-driven Experimental Research (POWDER) is part of the Platforms for Advanced Wireless Research (PAWR) program [1]. Development of POWDER is being done by the Flux research group at the University of Utah and takes place over a few years. The goal is to provide a testbed for researchers investigating next-generation wireless systems and mobile networks, including 5G. The testbed consists of base stations and endpoints placed around the University of Utah campus and Salt Lake City downtown. Each base station and endpoint has multiple software-defined radios and antennas. Some endpoints are mounted on mobile carriers such as buses. Users of POWDER create a project that is approved by the POWDER team and can then request resources. Resources are leased for a limited time to users, during which full administrative access is provided to the user.

Monitoring of the output of the software-defined radio itself is not sufficient to ensure proper use. Signals emitted by the radio pass through a power amplifier before being transmitted over-the-air. The power amplifier is based on solid-state semiconductor technology and can potentially produce spurious emissions. These include harmonics of the center frequency, intermodulation products, parasitics, and other frequency conversion products. These emissions have the potential to cause significant interference. The generation of digital samples for the SDR can be highly compute-intensive. Any kind of monitoring system must run independently on a dedicated system and monitor intentional transmissions as well as spurious emissions.

The monitor model shown in Figure 1 has been designed and implemented for each POWDER endpoint and base station. Users have access to an experimental SDR, whose output is amplified by a power amplifier (PA) and then output at an antenna. Directional couplers see a copy of the signal going to the antenna at a lower power level, but do not affect the signal sent by the user. To the user, the monitor is undetectable. The monitor receives from two channels simultaneously at R_1 and R_2 . These receive a combination of the output signal X and any incident signal Y originating in the external environment and picked up by the antenna. The directional couplers help separate the incident and transmitted signals due to their directionality.

Due to imperfect impedance matching between the PA, the couplers, and the antenna, reflections will occur in both directions. In particular, the transmitted signal will be reflected off of the antenna and incident signals will be reflected off of the PA. Any further reflections will not be considered as a simplification. The model for




Figure 1: A diagram of POWDER monitor system. The experimental SDR is used by experimenters to transmit out of the antenna or to receive signals from the antenna. A dedicated monitor SDR samples the spectrum from the directional coupler outputs.

the received signal then is,

$$R_1 = a_1X + a_2Y + a_3X + a_4Y \quad (1)$$

$$R_2 = a_5X + a_6Y + a_7X + a_8Y \quad (2)$$

The coefficients a_1, a_2, \dots, a_8 are complex values and represent the attenuation by couplers and reflection, and the phase offset due to the length of the transmission lines and also the reflections. Here a_1 represents the attenuation and phase offset applied to X as it reflects of the antenna and is received at R_1 . It is attenuated by approximately 23 dB. Of that 23 dB, 3 dB comes from reflection off the antenna, and 20 dB from going through the forward direction of the coupler. a_2 is the signal Y directly from the antenna and attenuated by 20 dB through the forward direction of the coupler. a_3 represents the attenuation applied to X , not reflected, going through the reverse direction of the coupler and is approximately 30 dB. Finally, a_4 is the attenuation of Y , reflected off the PA losing 3 dB and an additional 30 dB by going through the reverse direction of the coupler. This relationship is similar but not the same for R_2 .

A critical goal of the work presented in this paper is to isolate X . X is the signal transmitted by the user and will disappear if the user's experiment is shut down. If all the coefficients, a_1, a_2, \dots, a_8 , were known, it would be possible to solve for X since R_1 and R_2 are linear combinations of X and Y . The model can be simplified somewhat without loss of accuracy. Since a_1, a_2, \dots, a_8 are complex coefficients they can be combined.

$$R_1 = (a_1 + a_3)X + (a_2 + a_4)Y$$

$$R_1 = aX + cY$$

$$R_2 = (a_5 + a_7)X + (a_6 + a_8)Y$$

$$R_2 = bX + dY$$

Vector notation will be used going forward to simplify the expressions.

$$\begin{bmatrix} R_1 \\ R_2 \end{bmatrix} = \begin{bmatrix} a & c \\ b & d \end{bmatrix} \begin{bmatrix} X \\ Y \end{bmatrix} \quad (3)$$

To get X , take the inverse of the coefficient matrix and left-multiply by R_1 and R_2 .

To isolate X , the coefficient matrix must be estimated accurately. Estimation of the spectrum follows standard spectral estimation procedures after isolation. Once the spectrum has been estimated, powers are reported to the POWDER system by frequency in 100 kHz bins and compared against a database that records the maximum allowable power for that band. If power values exceed the database value, the experiment will be shutdown by the POWDER system.

To isolate X , the coefficient matrix must be estimated accurately. Estimation of the spectrum follows standard spectral estimation procedures after isolation. A fixed calibration procedure based on deconvolution is used to initially estimate coefficients. Deconvolution is used because the coefficients were found to be highly frequency dependent. Isolation of X was successful, but limited. It is anticipated that the estimates could be further improved via adaptive filter methods, independent components analysis, or other blind source separation methods. The relation of these methods to calibration-based source separation is discussed in section 2. Once the spectrum has been estimated, powers are reported to the POWDER system by frequency in 100 kHz bins and compared against a database that records the maximum allowable power for that band. If power values exceed the database value, the experimenter will be notified. If ongoing out-of-band emissions occur, the experiment can be shutdown by the POWDER system.

2 RELATED WORK

To realize an RF monitor, a receiver must be able to sample at very high rates or be able to quickly re-tune to successive center frequencies, sweeping across the spectrum to analyze large bands of spectrum. High rate receivers can be very costly [4]. Narrowband receivers require a re-tune of the LO, which can be time-consuming, and short duration signals of interest can be missed when the receiver is tuned to a different frequency band. SweepSense is a solution that uses low-cost narrowband receivers but, with modification to the LO hardware, incorporates a continuous tuning procedure to capture short-duration signals over a broad bandwidth [5]. Continuous tuning introduces distortion, which was solved in SweepSense by incorporating a self-calibration mechanism to remove distortion. Fast re-tunes are highly desirable, but since our work requires the deployment of many dozens of monitoring radios, hand-modification of the hardware was undesirable. In our signal model and our adversarial model, we can miss some short-duration signals without significant impact.

The National Telecommunications and Information Administration (NTIA) provided an overview of their Spectrum Monitoring Pilot program [3]. They design a system that is web-enabled and controlled. Their system, called SCOS, is open-source and available on Github. Spectrum sensors and measurements can be added to their system and can be controlled from a web user interface. The system is flexible and was initially considered as a framework for our work. However, it ultimately didn't fit the architecture POWDER currently uses.

Becker, Baset, et al., designed a spectrum monitoring system for real-time analysis. In addition to spectrum sensing, their system

was designed to classify signals. To do this, they limited their analysis to smaller frequency bands to find 802.11g, 802.11p, Bluetooth, and Zigbee signals. Their system also used USRPs, but instead of just UHD, they used GNU Radio. To meet real-time requirements, they split processing into a fast and slow component. The fast component, similar to our work, used energy detection on the frequency domain signal and then compared spectral parameters against a local database. The slow component used support vector machines to get more accurate classification results. A results merger system combined classification results and provided feedback to the fast component to improve capabilities over time. Our system also uses a lookup table for fast computation of results but does not currently use machine learning techniques.

Source separation is a critical component of our work. We need enough isolation between signal sources to identify energy in prohibited bands accurately. A substantial body of research exists for audio source separation, and much of it is applicable for radio signals, typically with some modifications. Audio signals are usually represented with real numbers, whereas RF signals almost always operate on quadrature I/Q samples, which are complex-valued.

Source separation techniques can be classified based on the number of sources and measurement channels. When the number of sources, n , is the same as the number of measurement channels m , independent components analysis (ICA) is often considered the state-of-the-art in source separation [6]. When the number of measurement channels is less than the number of sources, methods such as deep learning, non-negative matrix factorization (NMF) have been shown to be highly effective [2]. These methods all can take an initial estimate of the source separation matrix as a prior and improve it based on known stochastic properties of the signals.

3 THE POWDER RF MONITORING SYSTEM

One of the unique features of POWDER is that it is a so-called living laboratory. POWDER provides realistic environments, including dense, residential, and downtown regions. Deployments of POWDER radios here offer a unique opportunity to study next-generation wireless and mobile networking applications. Another essential feature is POWDER experimental profiles. These are programmatic interfaces to POWDER resources that include a request for hardware and virtual network links. They also specify disk images and data sets such that another user can copy the profile, repeat the experiment, and then extend it. As everything is programmable, compositions can be created that let novices and experts use POWDER effectively. A communications systems researcher can develop PHY layer protocols. Network research can develop 5G mobile network protocols without needing to re-invent the lower layers.

3.1 RF Equipment

POWDER wireless clusters are divided into two main types: base stations, which are high-power, high-performance systems, and fixed/mobile endpoints, which are lower power and could be mobile (e.g., when they are mounted to buses). On base stations, four wide-band USRPs are connected through an RF front-end and then to antennas. One antenna is banded, and the other is broadband. The front end provides filtering, amplification and provides a common point to monitor each of the radios. Multiple experiments can use

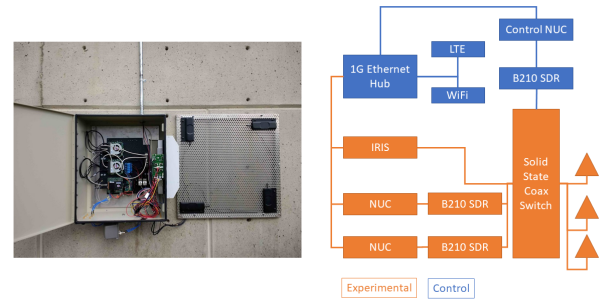


Figure 2: Hardware components of a POWDER fixed endpoint. Two USRP B210s are available to users. A third USRP B210 is used for monitoring. Intel NUC platforms provide local compute.

the same base station as long as they use separate radios and operate on separate frequency bands. On base stations, connections are made back to the data centers using a fiber backhaul. The other primary type of wireless node is a fixed/mobile endpoint. These nodes are shown in Figure 2. Endpoints are by far more numerous. These usually include two USRP B210 SDRs available for experimenters. Another RF front-end allows for monitoring, amplification, and filtering. As these endpoints may be mobile, two different backhaul connections are available. Campus WiFi (802.11) is extensive and high-rate and is the preferred connection type, but LTE is also available. Because data connections may sometimes be erratic and have lower rates than the base stations connected by fiber, local compute is available. An Intel NUC platform is connected to each of the B210s. Another NUC and B210 are used here, like in the base station, for monitoring and other platform utilities.

3.2 Monitoring Hardware

The main hardware components involved in monitoring include the control NUC, which is connected to the USRP B210 over USB. The B210 is connected over RF coaxial cable to the front end through the directional couplers and then out to an antenna. The NUC is an x86-64 based system with a quad-core Intel I7-8650 and 32 GB of DDR4 RAM. It supports USB3 and 802.11AC. The B210 supports USB 3.0 and performs all signal processing on a Xilinx FPGA. The RF front-end is custom built and has pairs of directional couplers as well as power and low-noise amplifiers.

USRP devices, generally can cover frequencies under 6 GHz. The B210 is able to sample at rates of up to 61 MSps. However, for monitoring, two simultaneous receive channels are needed. Using two channels limits the sample rate to 30.72 MSps. The lower limit of the B210 is 100 MHz. To cover the entire spectrum, the monitor must change center frequencies at intervals of 30.72 MHz, which would take 192 steps. Issues with filtering in the B210 lead us to choose smaller intervals and overlap the frequency bands slightly. Specifically, the automatically chosen analog filter bandwidth was too wide and did not remove enough out of band interference. Currently, it takes 238 steps to cover the spectrum. With these parameters, the B210 is the bottleneck in terms of speed performance.

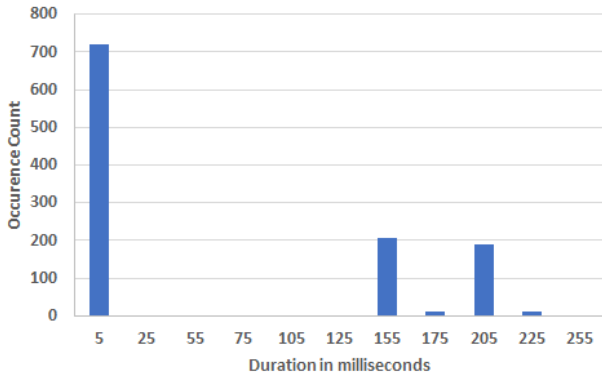


Figure 3: Re-tune elapsed time, histogram

The AD9361 provides the RF front end of the radio. The on-board PLL and ADC generally dictate the achievable rates and frequencies possible. When a new frequency is requested in software, a control command is sent to the B210. The FPGA translates this command to an AD9361 command, which is sent over an SPI bus to the AD9361. The chip commands the PLL to change frequencies. Some time is required for the PLL, which is driven by a feedback system, to lock on to the new frequency. If the chip detects that the new center frequency is over 100 MHz away from the last frequency, it initiates an on-board calibration. Tune time and calibration time, as the monitor swept across all 6 GHz, was a speed bottleneck. When frequency changes were small, the average time to re-tune was 4 ms. Every 100 MHz section, calibration took 100 ms on average. This latency leads to a total time spent on re-tunes per spectrum pass to be 6.7 seconds. This is shown in Figure 3.

One of the main components of the custom RF front end was the RF directional coupler system, as shown in Figure 4. The right image of Figure 4 labels the ports. These ports are bi-directional or symmetric, instead. When RF signals are transmitted through P1, they show up with minimal attenuation at the output port P2. The couplers are designed to pull a small amount of energy primarily from the P1-P2 signal. An ideal coupler would provide a lower energy copy of P1 at P3. Anything coming from P2, including reflections of P1, would not be present. As these devices are not ideal, some energy from P2, going in the reverse direction, is present at P3. We measured the directionality with a nearly matched antenna and estimated the isolation due to the couplers to be about 15 decibels.

3.3 Monitoring Software

The design of the software components of the monitoring system was undertaken with a few goals in mind. The monitor should be efficient and should be able to sweep over the available frequency as fast as possible. The system should be reliable, as many monitoring systems would be running concurrently and over long durations. Most importantly, the monitor system should be able to identify transmitted energy in out-of-band regions accurately. A general diagram of the software architecture is shown in Figure 5.

Nearly all software is implemented with Python 3. The most performance-critical pieces use C++ and C. The main libraries used




Figure 4: The directional coupler system. The left image shows a discrete four port coupler schematic. The center image shows four couplers on the same front end board. The right image shows a labeled diagram of each of the ports.

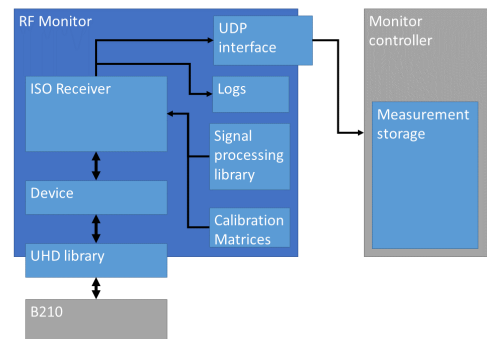


Figure 5: Software architecture of the monitor.

include NumPy for numerical processing, SciPy for a few signal processing utilities, UHD, and UHD python bindings for controlling and passing data to and from the B210, and XMLRPC. As shown in Figure 5, all monitor software was packaged as a python package for easy installation and updates across many clusters. The top-level monitor instantiates an isolation receiver object. This object, in turn, creates a device object. The device is capable of receiving samples from the B210 and handles all low-level UHD protocols. The device is mainly responsible for reliable B210 operation, too. The Isolation Receiver receives two streams of samples from the receive ports and applies an isolation algorithm to un-mix the transmitted and incident RF energy collected at the coupler. To change frequencies efficiently and without causing a sample stream hang, where samples stop coming from the device, we set the streamer to run continuously, even through frequency changes. This change substantially mitigated but didn't completely stop the stream hanging problem. We found that we did need to remove a certain number of samples after a tune because the data would be invalid before the PLL locks. Another supposed bug showed in the UHD driver software, where the software call to check PLL lock always immediately returned true. Our experiment showed conflicting results where signals that were known to be present at

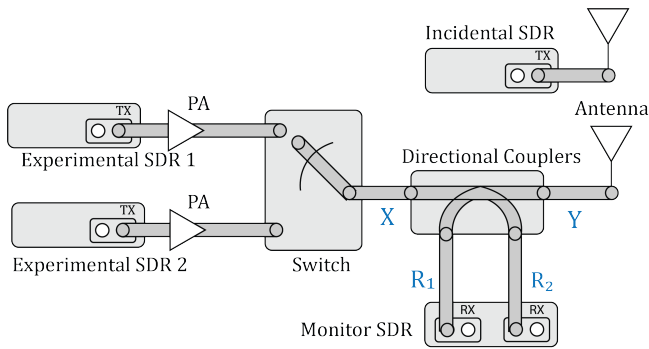


Figure 6: Calibration hardware is setup with multiple experimental SDRs under test. All compute nodes controlling the SDRs are connected via a TCP/IP network. Software running on the monitor compute node turns each experimental on and off, at the correct frequency, as needed. The monitor records spectrum measurements over each frequency band.

a particular frequency band would sometimes vanish. We implemented a time-dependent data dump mechanism to remove data based on how long the re-tune took. This solved the missing signal problem. Since re-tunes could take anywhere from 4 ms to 200 ms, and we couldn't afford to wait 200 ms each frequency change, a time-dependent mechanism was a reasonable design compromise.

A significant component of the monitoring isolation algorithm was measuring a known signal across the coupler ports and estimating the mixing matrix. We called this procedure monitoring calibration. Each mixing matrix is complex-valued and two by two shaped. We found that the mixing matrix was frequency-dependent, so we modeled mixing in the frequency domain and found a matrix for each discrete Fourier transform (DFT) coefficient. The derivation of this model is carried out later in this section, but for now, we use the results that an FFT size of 512 is appropriate. The set of measurement matrices thus included nearly 500,000 coefficients. To complete this efficiently, since it would need to be done for every radio used, the design shown in Figures 6, 7 was created.

When the TX server was started, it loaded a reference maximum length sequence (MLS) waveform and synchronized its local clock using the network time protocol (NTP). NTP was used because various calibration verifications done by the distributed application required a consistent ordering across systems. It used a cluster-specific configuration file to load networking configuration, port configuration, and other cluster-specific parameters. The TX server would not block while transmitting so that the client would immediately know that it had started broadcasting. A software lock was used on the PHY layer transmitter to prevent multiple calls to the server for broadcast. Each broadcast call was parameterized by a center frequency and a broadcast duration. The client would successively make calls to an appropriate transmitter, start recording when the lock was taken, and stop recording after a specified duration of time, but before the lock was released.

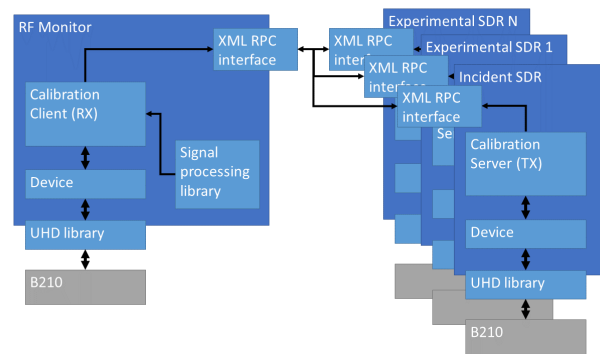


Figure 7: Calibration software consists of multiple distributed applications. The monitor compute node runs a calibration client, which provides control over the whole system. Each of the experimental nodes and the incident node run calibration servers. These provide RPC methods to the client that control a device object.

3.4 Source Mixing Model

Directional couplers provide some isolation on their own but extra signal demixing will need to take place to increase confidence in monitor output. Spectral measurements at the second RX port, R_2 , have increased energy contributions from the transmitter, X . Spectral measurements at the first RX port, R_1 , have increased energy contributions from the transmitter, Y . By comparing the spectrum, frequency bin by frequency bin, between R_1 and R_2 , a reasonable estimation of the transmitted energy can be made. However, due to imperfect impedance matching between the PA, the couplers, and the antenna, reflections occur in both directions.

An important empirical finding was that this coefficient matrix was fairly frequency-dependent. This is shown in Figure 8. Here a single position on the x-axis indicated one of the 230 center frequencies. Originally, we hypothesized that within 30 MHz, the frequency variability could be low, and a single coefficient would be appropriate. This graph shows that across the entire spectrum, the variability between coefficients 100 kHz apart could be substantial.

Instead of collecting one measurement per center frequency, a discrete Fourier transform (DFT) coefficient would be collected for each bin, and isolation would occur in the frequency domain instead of the time domain. Thus, the transfer function in the time domain is instead estimated.

The N -point DFT is defined as:

$$A(k) = \sum_{n=0}^{N-1} W_N^{kn} a_n = \mathfrak{F}[X(n)] \quad (4)$$

where

$$W_N = e^{-i\frac{2\pi}{N}} \quad (5)$$

and a_n is the input sequence to the DFT. The coefficients $A(k)$ represent the frequency components of a . $X(n)$ is the time domain sequence of samples collected, and \mathfrak{F} is the DFT operator. For each center frequency, f_c , and for each bin of the DFT, a coefficient matrix will be computed, for the entire set of spectral bins, F . N is set to




Figure 8: Initial estimates of the absolute value of a_1 over 6 GHz of spectrum. Each line represents an estimate at a different bin at the given center frequency. The x axis is an index into the list of the center frequencies. The droop around index 65 is between 2 and 3 GHz. The lowest blue lines are likely outliers. In general, significant variability exists as shown by the height of the band of values at each center frequency. Each point on the X axis represents 30.72 MHz of spectrum divided into 60 kHz bins

512. Importantly, the DFT operator is a linear operator, so the linear relations shown below hold.

$$\begin{bmatrix} \mathfrak{F}[R_1(n, f_c)] \\ \mathfrak{F}[R_2(n, f_c)] \end{bmatrix} = \begin{bmatrix} \alpha(k, f_c) & \gamma(k, f_c) \\ \beta(k, f_c) & \delta(k, f_c) \end{bmatrix} \begin{bmatrix} \mathfrak{F}[X(n, f_c)] \\ \mathfrak{F}[Y(n, f_c)] \end{bmatrix}, \quad (6)$$

which,

$$\begin{bmatrix} R_1(k) \\ R_2(k) \end{bmatrix} = \begin{bmatrix} \alpha(k) & \gamma(k) \\ \beta(k) & \delta(k) \end{bmatrix} \begin{bmatrix} X(k) \\ Y(k) \end{bmatrix} \quad (7)$$

For each RF port, 512 samples are collected in a sequence at 30.72 MSps. The sequence is windowed using a Hamming window, and then a 512-point DFT is computed. For each DFT coefficient, a matrix is estimated that can be used to transform the DFT coefficients found in R_1 and R_2 , to get the estimated DFT coefficient of X and Y . These can easily be converted to power spectral density measurements. Further, measurements are made of the actual power outputs, and a linear conversion is computed to get power in dBm from the computed dB measurements.

For general spectral measurements, we based parameters on FCC part 15 regulations, which place limits on emissions by commercial transmitters. Specifically, from section 15.35, an averaging detector with a period of 100 ms and bandwidth of 1 MHz is used for everything above 1 GHz. Below 1 GHz and supplementary to measurements above 1 GHz, a quasi-peak detector is used. This peak detector has a bandwidth of 100 kHz. 100 kHz was thus used as the minimum bin size for the monitor system. As 512 is the nearest power of two to 30.72 MHz divided by 100 kHz, 512 was used for the FFT size. At 30.72 MHz, each bin is sampled every 16 microseconds. We average bins over durations of up to 100 ms to get

high-SNR estimations of emissions. We plan to include maximum peak detections to align with the FCC commercial regulations more strictly.

3.5 Estimating Mixing Matrices

Robust, accurate estimation of the coefficients in A is needed to ensure good isolation. Monitor calibration is the procedure used to estimate these coefficients. A known signal is transmitted by the device being calibrated and is received by the monitor SDR at both ports. The monitor compares the received signal with the known signal and computes an estimate of A . The problem is made tractable by solving for the coefficients in two steps. First, X is set to zero by ensuring that only the incident transmitter is broadcasting (Y). Then Y is set to zero by ensuring that the incident transmitter is not broadcasting.

$$\begin{bmatrix} R_1(k) \\ R_2(k) \end{bmatrix} = \begin{bmatrix} \alpha(k) & \gamma(k) \\ \beta(k) & \delta(k) \end{bmatrix} \begin{bmatrix} X(k) \\ 0 \end{bmatrix} \quad (8)$$

$$\begin{bmatrix} R_1(k) \\ R_2(k) \end{bmatrix} = \begin{bmatrix} \alpha(k)X(k) \\ \beta(k)X(k) \end{bmatrix} \quad (9)$$

$$\begin{bmatrix} R_1(k)X^{-1}(k) \\ R_2(k)X^{-1}(k) \end{bmatrix} = \begin{bmatrix} \alpha(k) \\ \beta(k) \end{bmatrix} \quad (10)$$

$$\begin{bmatrix} R_1(k) \\ R_2(k) \end{bmatrix} = \begin{bmatrix} \alpha(k) & \gamma(k) \\ \beta(k) & \delta(k) \end{bmatrix} \begin{bmatrix} 0 \\ Y(k) \end{bmatrix} \quad (11)$$

$$\begin{bmatrix} R_1(k) \\ R_2(k) \end{bmatrix} = \begin{bmatrix} \gamma(k)Y(k) \\ \delta(k)Y(k) \end{bmatrix} \quad (12)$$

$$\begin{bmatrix} R_1(k)Y^{-1}(k) \\ R_2(k)Y^{-1}(k) \end{bmatrix} = \begin{bmatrix} \gamma(k) \\ \delta(k) \end{bmatrix} \quad (13)$$

By dividing the known signal DFT, X , from the received signals R_1 and R_2 , the operation is equivalent to deconvolution in the time domain with a long channel response. That is, if $x_1(n) \otimes x_2(n) = x_3(n)$, then $X_1(k)X_2(k) = X_3(k)$, so multiplication in the frequency-domain is convolution in the time domain. Special care must be given to the case where some frequency bins have values close to zero in the transmitted signal. The division presents a problem where the resulting computed coefficients may explode in value. No values in the DFT of X can be zero.

By using a pseudo-noise (PN) sequence, a known signal with a flat spectrum is generated. Specifically, a maximum length sequence (MLS) with the same length as the DFT is used. If an N -point DFT is used, the sequence repeats with a period N . The MLS sequence is exactly flat except for at the bin corresponding with frequency 0, which is often called the bias or DC value [7]. This value can be changed in the DFT representation by putting a non-zero value at the zero-bin or in the time domain by adding a constant number to the sequence to prevent the mean from being zero.

4 EXPERIMENTAL RESULTS

To test the performance of the isolation algorithm, measurements of the ability to remove incident energy were designed. In an ideal performance scenario, any incident signal energy is removed from the signal R_1 . Any transmitted signal energy is also removed from the signal R_2 . With just the directional couplers, a base level of performance is achieved. The directional couplers are able to suppress

energy from the undesired signals. For instance, with X set to zero (i.e. the transmitter is off), the spectra from Y will be present in both R_1 and R_2 but will have lower power on average. Increasing isolation by the algorithms described in the preceding sections should increase the difference in power levels between the desired signal and the undesired signal. This directly improves confidence in the detection scheme. By only reporting spectral measurements where the measurement is higher in the desired signal, spurious detections are minimized. Increasing the distance between the two spectrum measurements R_1 and R_2 , causes fewer false positives and false negatives.

Over the CBRS band, isolation between receiver and transmitter signals is estimated with an MLS sequence spectrum. As this signal has a flat spectrum, the average is taken at R_1 and R_2 . The difference is the pre-source demixing isolation. After calibration, the receiver measures R_1 and R_2 again but uses source demixing to increase the amount of isolation. This is the post-source demixing isolation. The average values for the CBRS band are shown in table 1. As seen in Figure 9, after isolation the difference between R_2 and R_1 is increased to 8 dB across the band.

Figure 10 is used to further illustrate the performance of the isolation algorithm. Figure 10 shows spectral measurements when both sides of the directional couplers are receiving signal energy. In this case, two sinc-like signals, which are rectangular functions in the frequency domain, are transmitted at different frequencies. X is a signal centered near 2.485 GHz and Y is a signal centered around 2.51 GHz. The bandwidth was chosen such that the signals did not overlap in the frequency domain. Under ideal performance, channel 1 would not measure any energy from Y . Additionally, channel 2 would not measure any energy from X . Isolation can be seen to be near 8 dB in both directions. The undesired signals are suppressed by about 8 dB from the desired signals. During monitoring, the Y signal will correctly not be reported.

The coupler provides some isolation, but isolation is frequency-dependent. By demixing the signal, isolation is more flat across the band and higher. Clipping was frequency-dependent, so hopefully, with a reasonable AGC solution, better isolation would be seen at other bands. An initial way the clipping problem was addressed was with frequency-specific gain settings. This solution is possible with the configuration files that are used during calibration as they site-specific and have a flexible structure. With more than 200 center frequencies and more than 100 sites, this solution needs some work to make automatic.

5 LIMITATIONS AND FUTURE WORK

This work provides an initial attempt and investigation into hardware-assisted source separation to provide monitoring of spurious emissions across a wide bandwidth with low-cost receivers. A variety of paths are open for continuation and improvement of this work. Of particular interest are increasing the tuning speed, increasing




Figure 9: A wideband MLS signal received on channel 1 and 2 then isolated by de-mixing. This figure shows spectral measurements when Y is the MLS and X is zero. Due to reflections, energy from Y shows up on chan 1, which is R_2 . Increasing the isolation causes this energy to be suppressed. Increasing isolation due to better calibration, better isolation algorithms, or better directional couplers will cause the difference between R_1 and R_2 to increase. The average distance between R_1 and R_2 when an MLS is used provides a good measure of performance.

the bandwidth, and increasing isolation. Some other limitations still exist with this monitor, but either has fixes planned or will be addressed in the future.

Currently, measured power output is relative and not in physical units of measurement. A reference is needed. On-going hardware measurements provide a mapping between monitor output and the real value in dBm. This mapping is being done as part of the deployment of each of the sites. Before deployment, a measurement is taken at various frequencies across the spectrum for each SDR. These are stored in a web server and are potentially accessible by the monitor application at run time.

At deployment sites where a powerful mobile communications transmitter is co-located with a POWDER transmitter, the ability to remove the external signals is limited. With the current amount of isolation, accurate identification of experimental emissions is difficult because the incident signal Y is much larger, relatively, than the experimental emissions meaning that demixing does not work, and effectively no isolation is present. In frequency-division multiple access systems, this is not a problem, because the experimental transmitter is not likely to cause interference at the mobile receivers due to having a lower transmission power. However, in time-division multiple access systems, an experimental transmitter can broadcast while the mobile receiver is receiving at the same frequency causing harmful interference. Currently, these sites are being evaluated on a case-by-case basis. RF front end notch filters may be used in these cases.

Table 1: Averaged isolation values in dB.

	isolation in dB experiment
Pre-source demixing	4.7438
Post-source demixing	8.2593

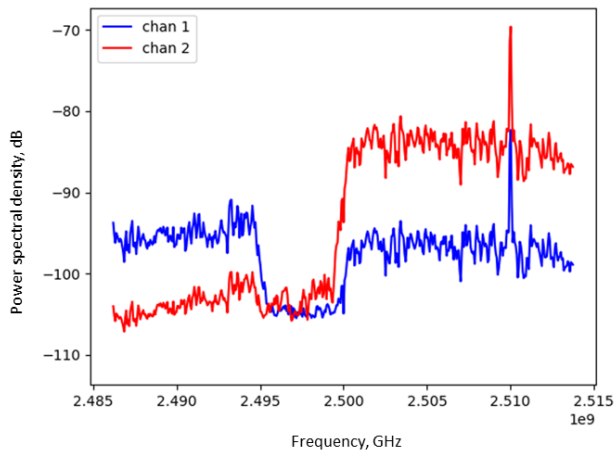


Figure 10: Two wideband MLS signals transmitted simultaneously by the experimental transmitter and the incident transmitter. The incident waveform is centered just below 2.485 GHz and the transmitted signal is centered at 2.510 GHz. The incident signal shows at a lower power on channel two and a higher power on channel one, clearly indicating that it is an external signal. The experimental signal is shown having higher power on channel one and lower power on channel two clearly indicating that it is an experimenters signal.

Transient signals are still a problem with this method. Tuning faster and at higher bandwidths show immediate improvements. Higher bandwidths are only possible with more capable receiver hardware. The B210-based system in this monitor operates at the maximum bandwidth possible for two channels. High-bandwidth receivers generally are substantially more costly. The tuning speed might be optimized by changing the properties of the AD9361 chip.

Improving isolation can be done in a variety of ways. Better isolation parameters in the couplers will help. Algorithmic solutions are also considered. Independent component analysis (ICA) is a promising method to increase isolation. For 2 signal sources and 2 measurement channels, it is state of the art for audio signal source separation. With only 1 measurement channel, deep learning methods are also promising. ICA uses an information-theoretic criterion to maximize the separation between two signals. Like the deconvolution method explained above, it estimates a mixing matrix. Deep learning methods use training data to learn the mixing matrix. With only a single channel, this is a much more difficult problem. Deep learning uses advances in neural network topologies, faster hardware, and large amounts of data to estimate complicated non-linear functions.

The POWDER platform enables cutting edge research into next-generation mobile networks and wireless communication systems. Next-generation systems will require larger bandwidths and more efficient usage of the spectrum. They will likely need new or improved sharing and multiple access mechanisms. As such, spectrum monitoring will be a vital component of these systems, particularly

for operators of communication infrastructure. The monitoring system discussed in this work is an attempt to improve upon state of the art for monitoring large spectrum bands with low-cost receivers. In this work, the objectives included a better understanding of the radio-frequency front ends used in POWDER and their role in adding spurious emissions that a user might not expect. This work also included a design of a spectrum monitoring system that could detect spurious emissions, and a better understanding of the issues of the spectrum monitoring system and its limitations.

ACKNOWLEDGEMENTS

This material is based upon work supported by the National Science Foundation under Grant Number 1827940. We acknowledge the support of the PAWR Project Office and the PAWR Industry consortium as well as our numerous partners at the University of Utah and in Salt Lake City.

REFERENCES

- [1] J. Breen, E. Eide, E. Lewis, D. Reading, A. Buffmire, M. Hibler, D. Maas, R. Ricci, J. Duerig, D. Johnson, A. Orange, D. Schurig, K. Dutt, S. K. Kasera, N. Patwari, L. B. Stoller, J. Van der Merwe, K. Webb, and G. Wong. Powder: Platform for Open Wireless Data-driven Experimental Research. In *ACM WiNTECH proceedings*, September 2020.
- [2] J.-T. Chien. *Source Separation and Machine Learning*. Academic Press, 2018.
- [3] M. Cotton, J. Wepman, J. Kub, S. Engelking, Y. Lo, H. Ottke, R. Kaiser, and D. Anderson. An Overview of the NTIA / NIST Spectrum Monitoring Pilot Program. 2015.
- [4] M. Ettus and M. Braun. The universal software radio peripheral (usrp) family of low-cost sdrs. *Opportunistic Spectrum Sharing and White Space Access: The Practical Reality*, pages 3–23, 2015.
- [5] Y. Guddeti, U. C. S. Diego, R. Subbaraman, I. I. T. Madras, M. Khazraee, A. Schulman, D. Bharadia, U. C. S. Diego, and I. Ndsi. SweepSense : Sensing 5 GHz in 5 Milliseconds with Low-cost Radios This paper is included in the Proceedings of the. *Proc. of NSDI*, 2019.
- [6] A. Hyvarinen, J. Karhunen, and E. Oja. Independent component analysis and blind source separation, 2001.
- [7] A. Milewski. Periodic sequences with optimal properties for channel estimation and fast start-up equalization. *IBM Journal of Research and Development*, 27(5):426–431, 1983.