

Searchlight: An accurate, sensitive, and fast radio frequency energy detection system

Richard Bell*, Kyle Watson[†], Tianyi Hu[†], Isamu Poy*, fred harris* and Dinesh Bharadia*

*Electrical and Computer Engineering, University of California, San Diego

[†]JASR Systems, San Diego, CA

Email: rbell@ucsd.edu, watson@jasr.systems, hu@jasr.systems, ipoy@ucsd.edu, fjharris@ucsd.edu, dineshb@ucsd.edu

Abstract—Detecting signals is often presented as a simple and solved problem. However, signal detection is typically not the only goal, as signal localization in time and frequency is often desired. This detection and localization process must also occur in the presence of non-ideal radio frequency (RF) receiver effects such as receiver non-flat noise floors, front-end inphase-quadrature (IQ) imbalance, local oscillator (LO) leakage, and in some cases power saturation. A robust signal detection system must overcome these hardware impairments and report only signals that correspond to true emitter devices. Additionally, some unspoken requirements take the form of detecting signals that are below the noise floor, very narrowband, very wideband, bursty, or frequency hopping. We present Searchlight, a signal detection system that solves these problems by using novel techniques that enable classical signal detection to work well using software defined radios (SDR). Performance results are presented for synthetically generated and over-the-air (OTA) data sets, which were obtained using commercially available hardware in a complicated signal scenario. Searchlight is shown to be an important enabler for true spectral monitoring platforms that desire to detect anomalous signals or apply machine learning (ML) algorithms to classify the various types of wireless activity in a given area.

Index Terms—Detection, Spectrum Analysis, Spectral Estimation, Polyphase Channelizer

I. INTRODUCTION

Detecting the presence of wireless energy is required by every wireless device in today's electronic world. For example, to transfer data using Wi-Fi, Bluetooth, LTE, 3G/4G/5G/6G, GPS, and any modern communication protocol, wireless access points must first detect the presence of a packet header, align to this header, and then decode the data. You might wonder, if this problem is so centrally important to our wireless lives, shouldn't it be solved by now? The answer to this question is that it depends on the amount of shared information between the transmitter and receiver

In 2020, the IARPA securing compartmented information with smart radio systems (SCISRS) program was announced [1]. The goal of SCISRS is to detect and characterize radio frequency (RF) anomalies between 100 MHz and 6 GHz to secure data in sensitive government spaces from wireless threats. Thus, it is not possible to rely on matched filters or other detectors that require prior knowledge to operate. An energy detector does not require prior knowledge, however, using an energy detector across 6 GHz of spectrum is difficult due to the varying nature of the signal protocols, modulations, and activity in the various bands. For example, the 2.4 GHz band is

very active and contains multiple signal types simultaneously. Alternatively, spectrum in the 3.5 GHz region is typically less active. The same parameterized detector in all bands will not yield good results. Fig. 1 shows a portion of the time/frequency (TF) plane with many different signals coexisting.

Searchlight operates across 6 GHz of spectrum to detect RF anomalies and is designed to meet the following requirements

- 1) Detect and localize (in TF) above and below noise floor.
- 2) Detect energy in any band from 100 MHz to 6 GHz.
- 3) Blind detection of signals.
- 4) Transform localized regions of energy from TF to time for further classification via machine learning (ML).

In Sect. I-A we present related works for signal detection. Sect. II describes the model of the incoming signal for a receive radio. Sect. III describes our solution to the detection problem, Searchlight. Sect. IV describes how the system is designed to be configured to allow detection across 6 GHz of spectrum. Sect. V presents Searchlight detection performance using synthetic and over-the-air data, and ablation studies to show how Searchlight algorithms improve upon standard techniques. Finally, Sect. VI presents the conclusion.

A. Related Work

Signal detection can be categorized into five main strategies; energy detection, matched filter detection, cyclostationary detection, covariance-based detection, and ML based detection [2]. Energy detectors have existed since the advent of radar during WWII. The theory of energy detection has also been well understood since the 1960's [3]. Research into energy detection systems today focus on improving the performance of energy detectors by finding optimal thresholds [4]–[6]. Overlooked is the process of defining how energy detection is implemented when signal properties like bandwidth and on-time are unknown. If detecting below noise floor signals is important, these issues matter. In addition, a discussion of how energy detection performance can depend on the signaling scenario in a particular band, whether it is highly congested or not, is overlooked. There is also very little information on the overall system requirements to implement an energy detection system that works with hardware on over-the-air data. Other improvements to energy detectors have been explored by improving the estimate of the power spectral density using techniques like wavelet transforms [7], [8]. While wavelets can improve the ability of a system to resolve signals in

certain regimes, it is not clear if they can be used to improve performance when signals of interest exist throughout the entire band, because at lower frequencies wavelet transforms produce poor time resolution, and at high frequencies they produce poorer frequency resolution [9]. Referring to Fig. 1, there are signals that exist at 0 MHz, -45 MHz, +45 MHz, and all frequencies in between. We cannot tile this space using wavelets without losing potentially valuable information at one TF, in favor of gaining information at another. Unless you can assign more value to a set of TF, it is not clear how one would use a wavelet transform to improve overall spectral resolution. Additionally, wavelet transforms do not support localized signal inversion as readily as the STFT, due to the non-uniform TF tiling of wavelet transforms. For example, it is not clear how the rectangles in Fig. 4 would be estimated and inverted using a scalogram.

Matched filters are known to be the optimal form of detection [10], but they require prior knowledge of the signal you are trying to detect. Cyclostationary based cycle detectors [11], [12] can outperform the energy detector in colored noise, however, they are difficult to use when prior knowledge of a signals cycle frequencies are unknown. They also tend to be computationally expensive for the general detection scenario and do not support inversion of the detected and localized signal to the time domain for further classification. Covariance based detectors use subspace decomposition to separate signal from noise by thresholding eigenvalues of the samples' covariance matrix [13], [14] which does not work well in low SNR scenarios since finding a proper threshold value becomes difficult. It also suffers from computational complexity related to computing eigenvalues, inhibiting its use in real-time systems. Deep learning based detection [15]–[17] is difficult to train for anomaly detection since anomalies by definition are unexpected signals. In addition, the domain gap between performance using synthetic data and performance over-the-air (OTA) is a serious problem, and remains an unsolved problem for ML in RF detection.

II. SYSTEM MODEL

With the goal of detecting individual signal energy at any center frequency, any bandwidth, and any on-time, that enters the antenna in the sub 6 GHz band, we will now present the model for the combined signal that makes our assumptions and task clear. The generalized passband signal model in a time varying multipath fading channel is

$$x_p(t) = A(t) \cos(2\pi f_c t + \phi(t)) * h_p(t, \tau), \quad (1)$$

where f_c is the carrier frequency, $\phi(t)$ is a time varying carrier phase, $A(t)$ is a time varying amplitude,

$$h_p(t, \tau) = \sum_{\ell=0}^{L-1} a_\ell(t) \delta(\tau - \tau_\ell(t)), \quad (2)$$

is the passband channel model for a non-stationary fading multipath channel with L paths, $\tau_\ell(t)$ and $a_\ell(t)$ are the time-varying delay and gain in the ℓ^{th} path respectively.

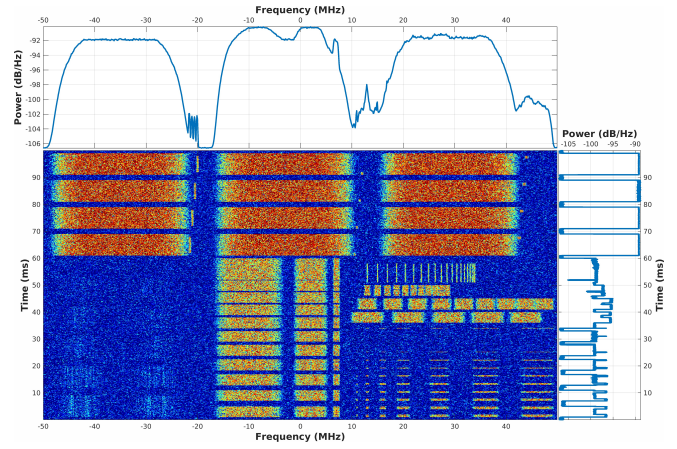


Fig. 1: A synthetic example of a complicated scenario in which many signals of varying powers, lengths, and bandwidths coexist in one batch of collected samples.

A receiver will downconvert the passband signal to baseband and sample the signal using an analog-to-digital converter (ADC). Suppose the receiver downconverts the received signal using a frequency f'_c , where $f_c - f'_c = \Delta f$ represents the carrier frequency offset (CFO). Now, suppose a sample period of T'_s is used by the ADC, where T_s represents the optimal sampling period. In addition, let α represent an offset in time from the optimal starting time for the sampling process, where $\alpha \leq T'_s$. Then $T_s - T'_s$ represents the sampling frequency offset (SFO) and α the sampling time offset (STO). The equivalent sampled complex baseband representation [18] of the received signal (1) incorporating these offsets can be expressed as

$$x(n) = \sum_{\ell=0}^{L-1} a_\ell A(nT'_s + \alpha - \tau_\ell) e^{j\phi(nT'_s + \alpha)} e^{-j2\pi\Delta f nT'_s} e^{-j2\pi\Delta f \alpha} e^{-j2\pi f_c \tau_\ell}, \quad (3)$$

where the dependence of a_ℓ and τ_ℓ on t has been dropped.

The detection system will collect N samples at any detection cycle. Within these N samples, there can exist M independent signals, denoted $x_1(n)$, $x_2(n)$, \dots , $x_M(n)$, each modeled by (3) with an appropriate set of parameters. An example spectrogram showing many signals within N samples of collected data is shown in Fig. 1. The final form of the received signal model becomes

$$y(n) = \sum_{m=1}^M x_m(n) + w(n), \quad (4)$$

where $w(n)$ is (potentially colored) Gaussian noise. The goal of Searchlight is to detect all $x_m(n)$ in the presence of $w(n)$ and all unknown parameters. For each detected $x_m(n)$, Searchlight estimates bandwidth, on-time, SNR , and reverts these regions to time domain to enable further classification.

III. PROPOSED SOLUTION: SEARCHLIGHT

Searchlight is a framework that solves the problem of detecting and localizing in the TF plane, all energies between 100

MHz and 6 GHz, while overcoming several challenges. The processing flow diagram for Searchlight is shown in Fig. 2. First, hardware impairments of radio receivers create “ghost” energies. These are copies of energy from one frequency to another that leads to false detects if unaccounted for. Searchlight implements a set of techniques to correct such impairments. Second, hardware radio receiver’s noise floor is non-flat and sometimes correlated. Accurate noise estimation is key to detecting signals that are below the noise floor, and Searchlight uses a novel technique to estimate the noise floor accurately. Third, radio receivers can sample signals at wideband, often as high as 100 MHz. Such wideband sampling creates the need for further localization in the TF plane. Searchlight uses a combination of convolution, power rate estimation, and sequential cancellation to achieve signal localization. Finally, the detected and localized region of the TF plane should be reverted back to band-limited complex baseband samples to enable further classification and/or parameter estimation. Searchlight achieves this using a polyphase synthesis filter.

Searchlight is designed to process streaming inphase-quadrature (IQ) samples in chunks of N , as supplied by a software defined radio. We will describe the various forms of processing these N samples go through shortly. If Searchlight detects M signals, after the processing is completed, then Searchlight’s output will consist of an IQ baseband representation of each of these signals within their localized region of the TF plane. In addition, metadata corresponding to start time, stop time, frequency low, frequency high, and SNR will be outputted for each detected IQ sequence. These are referenced with respect to the radio’s time reference and center frequency.

A. Hardware Impairment Removal

There are three major problems that the analog front-end of a radio can impart on a digital stream of samples that detrimentally affect energy detection performance; power saturation, IQ imbalance, and DC offset. If these effects are not accounted for, detection performance will suffer. The following describes how Searchlight solves the problem.

1) *Overcoming Front-End Power Saturation* : Power saturation occurs when the input signal power is too high. High power input can be the result of the gain of the receiver being too high, the placement of the antenna being too close to the emitter, or the emitter power being very high. Given any of these situations, voltage levels at the input to the ADC will largely map to the highest positive or negative bits at the output of the ADC, a phenomenon known as clipping. As a result of this non-linear transformation, energy becomes spread over wide regions of frequency, leading to incorrect localization and missed detects due to energy overlap.

Power saturation cannot be corrected after clipping has occurred at the ADC due to information loss. Our approach is to detect that saturation has occurred in the current sample stream, and prevent future occurrences in later streams. To do this, the chunk of raw (unprocessed) IQ is separated into a series of subchunks. Within each subchunk, the number of samples above a saturation threshold is counted, and a ratio of

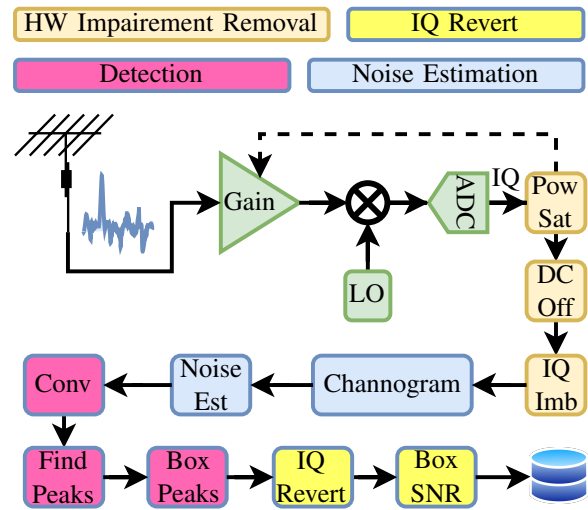


Fig. 2: The Searchlight processing steps from raw IQ input to boxed energy metadata at the output.

saturated samples to total samples is formed. The saturation threshold is a percentage of the full-scale range of the front-end ADC, and saturation is declared if the ratio rises above the saturation ratio threshold in a subchunk. The saturation level is defined as the average fraction of subchunks that were saturated within a full chunk. It was found that subchunks of length $10e3$, saturation threshold of 0.9, and saturation ratio threshold of $1e-3$ work well in practice. If one subchunk is saturated, the entire chunk is declared to be saturated, and Searchlight commands the radio to lower the front-end gain. Saturation detection and gain reduction iterate until power saturation is no longer detected.

2) *Overcoming IQ Imbalance*: One of the core functions of analog front-ends in a radio is to down convert signals from the passband to an intermediate or baseband frequency. As a result, low-cost receivers are enabled by reducing high-frequency routing and allows for lower-frequency ADCs to be used. Lowest-cost radios downconvert from passband to baseband before sampling, requiring separate analog paths to represent the inphase and quadrature portions of the signal. These independent analog paths impart different gains and cause a non-orthogonal phase between them. These imperfections cause IQ imbalance [19], and the net effect is to create attenuated copies of true emitter energy to the conjugate frequency. These energies do not exist outside the receiver, and thus are referred to as “ghost” energies. In the spectrogram, it is not possible to distinguish between “ghost” energies and real energies, because there is not enough information to make that decision once the phase has been stripped. Therefore, IQ imbalance must be corrected before the spectrogram is created.

To correct IQ imbalance after sampling, Searchlight measures the average gain on the I and Q parts of the complex signal and applies a correction to them so that they are equal. It then measures the average phase difference between the I and Q parts and applies a correction in the form of phase rotation so that they are orthogonal on average. These corrections

are estimated and applied to each IQ chunk, so that if IQ imbalance parameters change from chunk to chunk, they are always effectively removed.

3) *Overcoming DC Offset*: DC offset, which is narrowband energy centered at zero frequency, occurs for a variety of reasons inside radios. One major contributor is the local oscillator of a mixer mixing with itself via coupling between the input ports, creating what appears like a strong signal in the center of the spectrogram that will be detected and localized. Therefore, DC offset must be removed or accounted for, and there are two ways to achieve this. The first approach is to use a narrowband filter centered at DC to filter out the LO leakage contribution. The problem with this approach is that it filters all energy centered at DC, not just energy caused by DC offset. Consequently, contiguous energies that span DC appear like two distinct non-contiguous energies, and lead to false alarms. The second approach is to tune the center frequencies of the radio such that the spectrum of interest does not align with the center frequency of the radio, a technique known as offset tuning. For example, if the desired range of frequencies span 20 MHz from 200 MHz to 220 MHz, the radio can be tuned to a center frequency of 197 MHz using a sample rate of 50 MHz. In this way, all frequencies from 200 MHz to 220 MHz will not coincide with DC before ADC sampling. Once digitized, the stream can be downsampled and re-centered, producing better results at the expense of requiring higher receiver sample rate. Searchlight utilizes offset tuning to avoid DC offset.

B. Channogram and Noise Estimation

For Searchlight's detection and time and frequency localization requirements, it is necessary to transform the time domain IQ samples into a time and frequency representation to support natural searches over both degrees of freedom. The short-time Fourier transform (STFT) is ubiquitously used to support this type of search. The downside to using the STFT is that it distorts the frequency domain content through the choice of window used in the processing [20]. For example, the N -point FFT of a stream of N samples, sampled at f_s , produces a fundamental frequency resolution limit of f_s/N . As a result of windowing, the fundamental limit will be degraded by the choice of window and the dynamic range of the signals. The channogram, a super-resolution algorithm that overcomes this windowing degradation of the fundamental frequency resolution, was introduced in [21], and is used to replace the STFT, while maintaining the same computational complexity. Using previously corrected samples, Searchlight computes the channogram, a polyphase channelizer based estimate of the spectrogram, for a chunk of samples.

To determine if a set of samples contains signal or noise only, an estimate of the noise floor is required. Searchlight produces a noise floor estimate using the channogram, and sets a relative threshold to test if the total energy in the detector rises above or below it. The height of the threshold above the noise floor determines the probability of false alarm (P_{FA}) and probability of detection P_D . A robust and well known estimator for the noise floor is the median absolute deviation

(MAD) estimator [22]. MAD is robust because it rejects outliers in the data such as large narrowband line noise. However, it depends on the median, which means that if less than 50% of the samples are noise-only, the estimate will become poor. In regions of the spectrum that are congested, shown in Fig. 1, the signal sample count can become more than the noise-only count. When this happens the detector would underperform expectations or fail. To overcome this, Searchlight uses the minimum instead of median to estimate the noise floor. With the minimum, a noise floor estimate can be made when almost the entire chunk of IQ contains signal. To do this, Searchlight first performs upfront averaging to smooth the channogram, selects the minimum k samples, discards the lower m of this batch, and applies an empirical correction (that was found to work in practice) to transform the minimum to a mean value estimate. The minimum method works very well in practice and is robust across many signal scenarios for noise estimation. A performance comparison between the minimum and median based estimators is presented in Sect. V.

When samples are collected from hardware platforms, the assumption of a flat noise floor is often not met. In radios, there are analog filters that introduce correlations between samples, also known as noise coloring. If a minimum or median based noise estimator is applied to colored noise, it will produce very poor estimates. To account for this, Searchlight divides the band into subbands so that within each subband, the noise floor appears to be flat. The noise estimator is then applied to each subband and a parametric fit is made to combine the collection of estimates in each subband into one noise floor model. Using this noise floor model, the noise floor is subtracted from the channogram at each channel frequency and passed to the detector. Thus, non-flat coloring is effectively removed from the channogram in preparation for signal detection.

C. Detection

Energy detection has been studied since the advent of radar technology during WWII. Similarly, effective mathematical models have been published for almost as long [3]. What has not been addressed as thoroughly in the literature, is how to apply theory to maximize the probability of detection when no prior information is known. For example, the typical binary hypothesis test formulation of the energy detection problem assumes that when the signal exists, the energy detector length exactly matches the length of the signal and the signal occupies the detector's full band. If these conditions are not satisfied, the detector being used is sub-optimal and signals may be remain undetected. A useful relationship between the probability of detection (P_D), P_{FA} , detector length (N), and SNR at an energy detector input [23] is given by

$$P_d = Q\left(Q^{-1}(P_{fa}) - \sqrt{N}SNR\right), \quad (5)$$

where $Q(\cdot)$ is the Q-function for a normal distribution. While this equation is an approximation, it is applicable for most SNR and N experienced in practice. It also motivates the need to maximize the SNR within a detector of length N to maximize the probability of detection for a given probability

of false alarm. To maximize the SNR , non-overlapping noise, in both time and frequency, should be rejected. Of course, the problem of interest is to differentiate signal from noise, so how can noise be rejected before detecting the signal? This seems like a chicken and egg problem. Fortunately, there is a path forward that Searchlight uses to sidestep this problem that will be discussed next.

Using Parseval's theorem, we know how to compute the energy of a signal in either the time domain, $x(n)$, or the frequency domain, $X(k)$,

$$\sum_{n=0}^{N-1} |x(n)|^2 = \frac{1}{N} \sum_{k=0}^{N-1} |X(k)|^2. \quad (6)$$

In the TF plane, summing frequency samples across multiple rows is an extension of Parseval's theorem. The process of summing entire rows together to compute the energy can be modified to summing those regions that lie within a rectangle in the plane, producing an estimate of the energy over a limited portion of time and frequency. A visualization is shown in Fig. 3 where three line plots and the TF plot correspond to the same underlying time domain IQ sequence. Limiting the vertical height of the rectangles in the time/freq plane corresponds directly to the horizontal width of the rectangles in the time domain. Limiting the horizontal width of the rectangles in the time/freq plane cannot be visualized in the time domain, but is equivalent to a passband filter.

Detector length translates to the total number of TF bins (or pixels) contained within the rectangle, also known as the time-bandwidth product. Therefore, if a large number of convolution kernels is chosen with varying time-bandwidth products, detector length is searched over in addition to time and frequency. If every possible rectangle that fits within the channogram is exhaustively searched, every possible energy detector that can be applied to the time domain samples is also explored. Searchlight uses this relationship to apply a large set of detectors to every chunk of IQ, making undetected signals less likely. Using this technique, Searchlight is able to detect below noise floor as well as above noise floor signal energy.

When the time-bandwidth product of a convolution kernel gets larger, non-contiguous energies can start bleeding into one another. If the kernel simultaneously overlaps them, it will lead to imprecise or incorrect detection and localization. To combat this, Searchlight uses a successive detection and cancellation procedure to remove energies from the channogram as they are detected. The search through convolution kernels always begins with small and ends with large time-bandwidth products, ensuring that high power energies are found first and canceled. The peak value of the convolution output is found and the boxing process begins. From here, we move left/right/up/down from the peak, while computing the rate of change of power. When the power rate change surpasses a threshold, a box edge is declared. After four edges are estimated, pixels within the box are zeroed out, and a new round of convolution is made using the same kernel over the new channogram until no more peaks are found above the

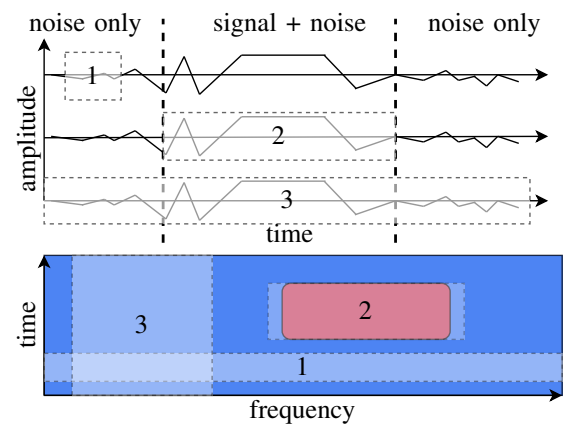


Fig. 3: Visual representation of rectangles in the time domain to rectangles in the TF plane.

detection threshold. At this point, the next convolution kernel is repeatedly selected until all kernels are exhausted.

D. IQ Revert and Parameter Estimation

When all energies have been detected and localized, there are three remaining functions that Searchlight is designed to provide; overall rate reduction, reverting the TF localized region to the time domain, and parameter estimation for each localized energy. Rate reduction is important because it frees resources on the host computer for other operations such as signal classification, or lowers network congestion if Searchlight results are sent to a centralized processing node. Reverting TF samples to time domain enables classification algorithms, requiring time domain data, to be used. Finally, parameter estimation for useful/required information can assist in decision making for further classification.

Searchlight transforms each energy localized region of the channogram back into IQ time domain samples using a synthesis polyphase channelizer as described in [21]. Such information is useful for algorithms that require time domain samples for further classification. For example, a strip spectral correlation analyzer [24], a form of cyclostationary estimator, requires time domain samples to estimate the spectral correlation density (SCD). The SCD can then be passed through a neural network to aid in further classification of each detected and localized energy. Some examples of additional labels might include the modulation, the protocol, and/or the modality (single carrier, multi-carrier, dsss, etc.) for each localized energy. This inversion of localized regions of the channogram using a synthesis polyphase channelizer is much more computationally efficient than accessing a stored version of the original (high sample rate) IQ and filtering this in separate stages. The work has already been done to transform to the TF plane and determine the occupied bandwidth of the signal, do not throw this away only to reproduce the processing using a classical filtering method or approach.

Searchlight is capable of estimating parameters of detected and localized energy. Currently, it estimates the SNR of each

detected and localized energy by computing the total energy within the rectangle, subtracting the estimated noise power using the noise floor estimate in that region, and dividing this difference by the noise power estimate in that region

$$SNR = \frac{\hat{P}_{\text{box}} - \hat{P}_n}{\hat{P}_n}. \quad (7)$$

This is useful when trying to differentiate between energies originating from different radios, or as another dimension for clustering algorithms applied to all boxes. Additional parameter estimates can easily be added to Searchlight.

IV. REAL-TIME RECONFIGURABILITY

Not all regions of spectrum are created equal. The spectral content below 1 GHz takes on very different forms from the spectrum at 2.4 GHz and then again from the spectrum at 5.7 GHz. Not only is the signal content different, but the noise floor varies between these regions as well. Assuming that one Searchlight configuration will work equally well at every part of the spectrum from 100 MHz to 6 GHz will lead to failure. Therefore, it supports real-time reconfiguration as the center frequency changes. For each center frequency, a configuration file is defined and loaded before processing.

We have identified the following parameters whose value should depend on the center frequency of the collection

- 1) Assumed error in noise floor estimate - this value is added to the detection threshold for fine tuning detection performance. In regions of dirty RF, it should be made larger to maintain the same level of P_{FA} .
- 2) Boxing power rate threshold - the threshold that the power rate change is compared to when determining the edges of a box. Some regions of the spectrum contain signals with faster power rate changes and some with slower power rate changes.
- 3) Upfront averaging - the amount of smoothing that is applied to the channogram before noise estimation and detection takes place. A useful parameter that helps produce reasonable boxes in complicated scenarios that might otherwise be difficult to handle, even for a human.

Data was processed using Searchlight in 59 subbands, ranging from 150 MHz to 5.95 GHz in steps of 100 MHz. In regions below 1 GHz, many signals are always on. The spectrum in these regions tends to be dirty, meaning there is a lot of variation in energy that is difficult to attribute to a signal or to noise. The assumed error in noise floor is raised to 0.6 dB in the lower ends of this region, slowly lowering this value to 0.2 dB as the center frequency increases. Averaging in time is increased to 2 ms to avoid producing multiple boxes, and the power rate threshold is increased so that larger changes are required to declare an edge. These settings produced better overall boxing results for this region of frequencies.

For center frequencies between 3 GHz and 5 GHz, there is mostly unused spectrum, in the San Diego area around UCSD. Here, the assumed error in noise floor is reduced to 0.1 dB, and the amount of averaging to 1 ms. The boxing power rate threshold is also reduced so that box edges are declared sooner.

This is possible because the spectrum is clean and the energy density is low, so there is less interference between energies causing degradation in boxing below 1 GHz.

V. RESULTS

To evaluate detection systems that make claims like Searchlight, it is important to test systems over-the-air (OTA) and against a number of signal scenarios that stress the system. This ensures that the system is not independently tuned for one scenario, to the detriment of all others. With this in mind, the primary test scenario used to produce Searchlight performance results is shown in Fig. 4. The SNR for all signals shown in the figure is set to 10 dB so they can be visualized clearly. During a test, the SNR would be lowered accordingly. The rectangles drawn over the channogram correspond to Searchlight's detection and localization results, and numbered rectangles correspond to various failure cases made by Searchlight.

There are five signal scenario zones defined in Fig. 4. The first scenario on the bottom left, ranging from 0 ms to 50 ms and -50 to -25 MHz for a total of 10 rectangles, tests below noise floor performance. These energies are modulated using direct sequence spread spectrum (DSSS) modulation, with SNR starting at 0 dB on the bottom left, falling to -18 dB SNR on the top right. The second scenario to the right of the first, ranging from 0 ms to 60 ms and -25 MHz to 10 MHz, tests time resolution performance. These are three different bandwidths of quadrature amplitude modulated (QAM) energies whose spacing in time begins at 3 ms and decreases to 4 μ s from bottom to top. The next scenario at the bottom right, ranging from 0 ms to 35 ms and 10 MHz to 50 MHz, tests varying time-bandwidth product energies. These are QAM modulated signals ranging from a time-bandwidth product of 4500 at the bottom right to 1 at the top left. The region above this, from 35 ms to 60 ms and 10 MHz to 50 MHz, tests frequency resolution performance. The signals are spaced 3 MHz to 100 kHz, from left to right. The final scenario is found in the upper third of the channogram, ranging from 60 ms to 100 ms and -50 MHz to 50 MHz. This region tests for snugglers, which consist of a pair of signals, the overt signal and the snuggler. The overt is used as cover by the snuggler to hide against. Snugglers in this region begin 2 MHz away from their overt and end 0.5 MHz from their overt, from top to bottom. The snugglers on time and relative location change from left to right in this region. Searchlight was tested using both synthetic (made purely from software) and OTA (signals that underwent a physical transmitter, wireless channel, and receiver) forms of the combined scenario shown in Fig. 4. Searchlight performance is compared using this scenario to determine the amount of degradation one can expect when transitioning from synthetic to OTA data.

Performance results of synthetic data for the combined scenario of Fig. 4, are shown in Table I. The cause of a lower than expected P_D can be explained by studying the failure regions in Fig. 4. Failure region 1, shown as the white box with a circled number 1 beside it, corresponds to Searchlight making two boxes when it should have only made one. This

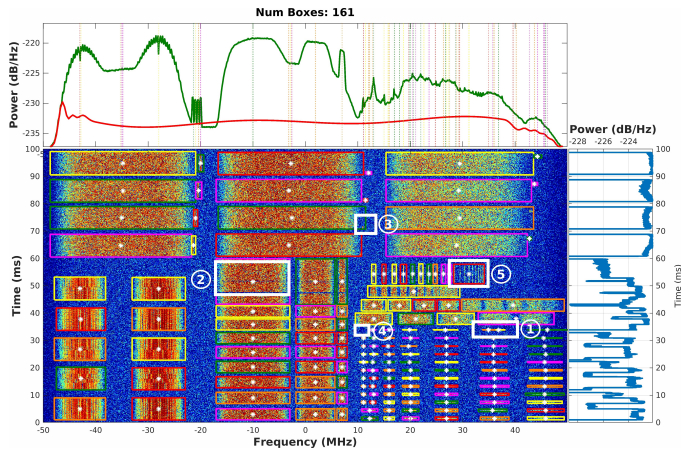


Fig. 4: Detected energies within a channogram. Numbered regions highlight failure modes that degrade TPR and FPR.

is referred to as a multibox failure, which tends to cause P_{FA} to increase, because the second box will not correspond to a true signal, given that the first box will have been associated with the truth already. Failure regions 2, 3, and 5 occur when Searchlight is not able to resolve multiple energies as distinct, and instead draws one rectangle containing the entire set. P_D drops as a result, because the one box will be associated with only one of the enclosed truth signals, leaving the remainder to be unassociated. This is a major failure source causing P_D to be lower than expected. The P_D would be 95% or higher if these failure scenarios were removed. Region 4 shows not boxing an energy that is present with a time-bandwidth product of 1, meaning the signal consists of one time domain sample.

Intersection over union (IoU), a measure of how well the predicted rectangle overlaps the truth rectangle, can be a misleading result for RF localization because the truth box is defined using the 3 dB definition of bandwidth. This can result in boxes that do not appear to capture a majority of the signal energy on a log plot in the TF plane. Searchlight was tuned to predict rectangles that capture 99% of visible energy, leading to rectangles that tend to be larger than truth assuming a 3 dB definition of bandwidth, and result in intersection over union that deviate from close overlap. Searchlight allows for rectangle prediction tuning if it is important that predicted boxes match a particular truth bandwidth definition.

The frequency resolution, Δf , and time resolution, Δt , in Table I correspond to the minimum resolutions before energies are combined into one box. While it is certainly possible to tune Searchlight for finer resolution results, it is more difficult to do this while maintaining good performance across all scenarios and bands. When SNR decreases, these resolutions cannot be measured because the boxes are not always detected. For these results, Searchlight used 1024 channels at a sample rate of 100 MHz, corresponding to a fundamental frequency resolution of 98 kHz and a fundamental time resolution of 4 μ s. These limits are further degraded by the amount of upfront averaging, four pixels in this case, resulting in resolutions

TABLE I

Detection performance per SNR for synthetic data.

SNR (dB)	P_D %	P_{FA} %	IoU %	Δf (MHz)	Δt (ms)
10	66.4	10.7	60.02	0.8	0.6
-2	66.4	3.0	63.13	0.8	0.6
-4	64.3	1.5	62.86	1	0.6
-6	55.6	1.5	60.79	1.2	0.9
-8	31.2	1.2	60.01	—	—
-10	8.9	0.7	62.31	—	—
-12	1.2	0.2	45.59	—	—

TABLE II

Detection performance per SNR for OTA data

SNR (dB)	P_D %	P_{FA} %	IoU %	Δf (MHz)	Δt (ms)
10	54.2	12.8	55.3	0.8	0.6
-2	38.4	2.4	58.0	0.8	0.6
-4	36.4	1.7	56.0	1	0.6
-6	15.9	0.5	61.9	1.2	0.9
-8	4.3	0.09	66.8	—	—
-10	0.27	0.29	26.8	—	—
-12	0	0.50	0	—	—

of 392 MHz and 16 μ s, respectively. Finally, the power rate threshold used when estimating box edges will degrade resolution further. Degradation depends on a signal's power profile, if the signal's power changes fast, less will occur.

Table II shows performance metrics against the scenario shown in Fig. 1 OTA. There is degradation in performance for the three main metrics. The degradation is attributed to two major causes: 1. the wireless channel introducing effects that cause multibox errors to occur and 2. how SNR is calibrated between synthetic and OTA scenarios. Multipath fading, for example, can cause one contiguous energy to appear as two after passing through the channel. It is also difficult to tune an OTA system to produce specific SNRs at the receiver. The solution is to tune the receiver gain until a received test signal matches the height above noise floor for corresponding synthetic signals at a given SNR. Accuracy was within 1-2 dB of true SNR on average, and IoU performance may be further degraded by non-flat noise floor using OTA data.

To demonstrate improvement's in detection performance, an ablation study is performed by replacing one algorithm at a time with different, standard approaches. Changes in performance are plotted in Fig. 5. For this study, three specific algorithms were tested; iteration through convolution kernel size and shape, noise floor estimator, and the sequential detection and cancellation algorithm. The line labeled "Searchlight" is Searchlight's performance using the algorithms described in Sect. III. "Fixed Kernel" refers to disabling the search over convolution kernel size, done to utilize the optimal detector for each energy, using only one fixed kernel instead. The figure shows that this leads to a decrease in P_D at each SNR. "K-means" refers to replacing the successive detection and cancellation algorithm with a K-means clustering algorithm. All points above the energy detection threshold are used to form clusters. K-means is run for a sequentially increasing number




Fig. 5: Ablation study performance comparison to Searchlight.

of hypothesized means until the total error stops dramatically decreasing. The K-means method produces poorer P_D and P_{fa} . “Noise - Median” refers to using a more standard median-based noise estimator. While P_D decreases when using the median, the true advantage of the minimum estimator appears when signals occupy more than 50% of a chunk’s pixels. When this occurs, the median-based estimator fails, causing a collapse in P_D (not shown in the plot).

VI. CONCLUSION

This paper presents Searchlight, a novel detection framework that enables detection and localization, in the time/frequency plane, of energies above and below the noise floor from 100 MHz to 6 GHz. This is achieved with a combination of algorithms that allow Searchlight to operate in the presence of hardware impairments, and to localize any signal energy. Searchlight achieves a P_D of 36% with a false alarm rate of 1.7% using over-the-air data at -4 dB SNR in a complicated scenario with over 200 energies. This includes smugglers, close frequency, close time, small time-bandwidth product, and below noise floor signals. In simpler scenarios where energies are well-spaced, the P_D increases dramatically. Searchlight is a dependable detection system that can be used with software defined radios to secure wireless spaces and enable spectrum sharing. In future work, we will further improve the detection performance in complicated scenarios by introducing improvements to the energy detection process.

VII. ACKNOWLEDGEMENTS

This section is based upon work supported in part by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via [2021-2106240007]. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

REFERENCES

- [1] “Securing compartmented information with smart radio systems (SCISRS),” <https://www.iarpa.gov/research-programs/scisrs>, accessed: 2022-12-07.
- [2] Y. Arjoun and N. Kaabouch, “A comprehensive survey on spectrum sensing in cognitive radio networks: Recent advances, new challenges, and future research directions,” *Sensors*, vol. 19, no. 1, p. 126, 2019.
- [3] H. Urkowitz, “Energy detection of unknown deterministic signals,” *Proceedings of the IEEE*, vol. 55, no. 4, pp. 523–531, 1967.
- [4] S. Atapattu, C. Tellambura, and H. Jiang, “Energy detection based cooperative spectrum sensing in cognitive radio networks,” *IEEE Transactions on Wireless Communications*, vol. 10, no. 4, pp. 1232–1241, 2011.
- [5] I. Develi *et al.*, “Spectrum sensing in cognitive radio networks: threshold optimization and analysis,” *EURASIP Journal on Wireless Communications and Networking*, vol. 2020, no. 1, pp. 1–19, 2020.
- [6] Smriti and C. Chhagan, “Double threshold-based energy detection spectrum sensing scheme by considering the sensing history in confusion region,” in *2018 5th International Conference on Signal Processing and Integrated Networks (SPIN)*, 2018, pp. 518–521.
- [7] S. Kim, Y. Yoon, H. Jeon, M. Kim, and H. Lee, “Selective discrete wavelet packet transform-based energy detector for cognitive radios,” in *MILCOM 2008 - 2008 IEEE Military Communications Conference*, 2008, pp. 1–6.
- [8] A. Kumar, S. Saha, and R. Bhattacharya, “Wavelet transform based novel edge detection algorithms for wideband spectrum sensing in crns,” *AEU-International Journal of Electronics and Communications*, vol. 84, pp. 100–110, 2018.
- [9] S. Mallat, *A wavelet tour of signal processing*. Elsevier, 1999.
- [10] D. Bhargavi and C. R. Murthy, “Performance comparison of energy, matched-filter and cyclostationarity-based spectrum sensing,” in *2010 IEEE 11th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*, 2010, pp. 1–5.
- [11] W. Gardner and C. Spooner, “Signal interception: performance advantages of cyclic-feature detectors,” *IEEE Transactions on Communications*, vol. 40, no. 1, pp. 149–159, 1992.
- [12] C. M. Spooner, A. N. Mody, J. Chuang, and M. P. Anthony, “Tunnelized cyclostationary signal processing: A novel approach to low-energy spectrum sensing,” in *MILCOM 2013 - 2013 IEEE Military Communications Conference*, 2013, pp. 811–816.
- [13] Y. Zeng, C. L. Koh, and Y.-C. Liang, “Maximum eigenvalue detection: Theory and application,” in *2008 IEEE international conference on communications*. IEEE, 2008, pp. 4160–4164.
- [14] K. S. Kumar, R. Saravanan, and R. Muthaiah, “Cognitive radio spectrum sensing algorithms based on eigenvalue and covariance methods,” *Int. J. Eng. Technol.*, vol. 5, no. 2, pp. 385–395, 2013.
- [15] A. Vagollari, V. Schram, W. Wicke, M. Hirschbeck, and W. Gerstaecker, “Joint detection and classification of rf signals using deep learning,” in *2021 IEEE 93rd Vehicular Technology Conference (VTC2021-Spring)*, 2021, pp. 1–7.
- [16] X. Zha, H. Peng, X. Qin, G. Li, and S. Yang, “A deep learning framework for signal detection and modulation classification,” *Sensors*, vol. 19, no. 18, p. 4042, 2019.
- [17] Y. Mo, J. Huang, and G. Qian, “Deep learning approach to uav detection and classification by using compressively sensed rf signal,” *Sensors*, vol. 22, no. 8, p. 3072, 2022.
- [18] D. Tse and P. Viswanath, *Fundamentals of wireless communication*. Cambridge university press, 2005.
- [19] M. Valkama, M. Renfors, and V. Koivunen, “Advanced methods for i/q imbalance compensation in communication receivers,” *IEEE Transactions on Signal Processing*, vol. 49, no. 10, pp. 2335–2344, 2001.
- [20] F. Harris, “On the use of windows for harmonic analysis with the discrete fourier transform,” *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51–83, 1978.
- [21] R. Bell, f. j. harris, P. Gerstoft, and D. Bharadia, “High resolution spectral analysis and signal segregation using the polyphase channelizer,” in *2022 56th Asilomar Conference on Signals, Systems, and Computers*, 2022, pp. 519–526.
- [22] C. Leys, C. Ley, O. Klein, P. Bernard, and L. Licata, “Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median,” *Journal of experimental social psychology*, vol. 49, no. 4, pp. 764–766, 2013.
- [23] S. M. Kay, *Fundamentals of statistical signal processing, Volume II, Detection Theory*. Prentice Hall New Jersey, 1998.
- [24] W. Brown and H. Loomis, “Digital implementations of spectral correlation analyzers,” *IEEE Transactions on Signal Processing*, vol. 41, no. 2, pp. 703–720, 1993.