Towards Physical Layer Identification of Cognitive Radio Devices

Seth Andrews*, Ryan M. Gerdes*, Ming Li†
* Virginia Tech
†The University of Arizona

***Abstract*—Increasing demand has led to wireless spectrum shortages, and many parts of the existing spectrum are heavily used. Dynamic spectrum access (DSA) has been proposed to allow cognitive radio networks to use existing spectrum more efficiently. It will allow secondary users to transmit on already allocated spectrum on a non-interference basis. Cognitive radios are able to change bandwidth and other transmission characteristics to take advantage of this spectrum. To enforce spectrum access rules it is necessary to uniquely identify all devices on the network. Manufacturing variation cause every device to have minute differences. Physical layer identification (also called device fingerprinting) techniques allow identification of devices based on small but unique variation due to these imperfections. Fingerprinting is very sensitive to any changes in the capture setup or device's environment. The changes in bandwidth that would occur in a DSA system cause device fingerprinting to fail. In this paper, we extend current device identification methods to include identification of devices with changing bandwidth. Experimental results are demonstrated on a collection of over 50 transmitters, with a significant improvement over current methods.**

I. Introduction

As wireless transceivers find an ever increasing number of applications the efficient allocation of spectrum has become an important topic. Many portions of the radio spectrum are unregulated, leading to interference among competing users. Licenses to other portions of the spectrum have been sold. This leads to large chunks of spectrum that are unused at various times, in different geographical locations, or in some subset of the spectrum, when the licensed user (primary user (PU)) does not use it. As demand for wireless spectrum increases, more efficient spectrum allocation is needed. DSA networks will allow cognitive radios to dynamically allocate unused spectrum, and to assign it among various users in the network.

Allowing transceivers to operate in this manner introduces several security concerns. Selfish users can transmit outside their allowed bandwidth, power levels, or time intervals. This could cause interference to others in the DSA network, and to licensed PUs. For these reasons it is useful to identify and track cognitive radio transmitters to enforce spectrum access rules. It is desirable that no modifications are made to the PU's transmission so that compatibility is maintained with legacy systems with spectrum licenses. While typical security identifiers, such as MAC addresses or encryption keys, could serve this purpose, they can be stolen (in the case of encryption keys) or easily forged (for MAC addresses), and introduce overhead to the system.

Physical layer identification (PLI) allows identifying devices based on what they are, rather than something they have [1], [2]. It is often referred to as device fingerprinting, in reference to traditional fingerprinting techniques that identify individuals based on unique biometric markers. PLI relies on identifying small differences in each device's output that occur due to age, circuitry, and manufacturing variation in otherwise identical models. Authorized users are "enrolled" in the system by obtaining examples of their signal. When a device asserts an identity, it is compared with the enrolled signals for that device. When the comparison is sufficiently close the device is accepted. In this way authorized devices can be identified with no additional overhead to the system.

However, the variations in a cognitive radio's signal due to changing bandwidth is much larger than the variation unique to each device. This causes current fingerprinting techniques to fail. We propose a new fingerprinting method which allows identifying devices with changing bandwidth, and which could be extended to other transmission parameters.

A. Cognitive radio

Software defined radios (SDRs) are able to dynamically reconfigure transmission parameters such as carrier frequency, bandwidth, transmit power, and even modulation waveform. Software defined radios and traditional transmitters share many of the same components, including digital to analog converters, mixers, and pulse shaping filters. However, in a cognitive radio system the hardware is controlled by higher level protocol layers. This allows dynamically adjusting transmission parameters (such as changing bandwidth or modulation type) in response to changing channel conditions and application demands [3].

The dynamic nature of cognitive radios makes fingerprinting challenging. PLI methods are very sensitive to the training data enrolled for each device. The enrolled data needs to be a portion of the signal that is the same in all transmissions by the device. Any change in the signal received reduces the effectiveness of the system: even using different antennas significantly degrades performance [4]. A device in a DSA system may change its transmission parameters to use different modulation types and transmit at several carrier frequencies and bandwidths. Due to the dynamic nature of the system, it may not be possible to know all parameters beforehand. Changes in bandwidth will cause standard fingerprinting techniques to fail.

B. Contributions

We propose what is, to the best of our knowledge, the first fingerprinting system for devices transmitting at multiple

bandwidths and demonstrate experimental results on a cognitive radio system using eleven bandwidths. Currently, any large changes in a device's bandwidth requires an entirely new set of reference data to re-identify the device. This is impractical, and increases the complexity of PLI systems. We propose a transfer learning method for the identification of devices with changing transmission parameters and limited reference data. A small subset of devices are enrolled at each bandwidth, and used to describe intra-bandwidth relationships. A method for choosing how to enroll devices to minimize the number of device and bandwidth combinations is given.

Experimental results validate the proposed method using data gathered from 50 transmitters operating at eleven different bandwidths. Several variations on the method are compared, and all perform significantly better than current techniques. This is compared with performance using another transfer learning technique, without any transformation, and with the ideal performance when the device has reference signals at the transmission bandwidth. We also point to future directions for transfer learning applied to radio fingerprinting, including varying carrier frequency and modulation types.

C. Paper structure

The paper is arranged as follows: Section II covers several preliminary points including the basic steps for PLI, an introduction to transfer learning techniques, and a description of feature extraction. In Section III an overview of physical layer identification (PLI) for cognitive radios is given along with possible attacks on this system. In Section IV we describe a robust method for fingerprinting cognitive radio devices with changing transmission parameters. This is followed by a description of the experimental setup and results in Section V. We end with a summary of related work on securing cognitive radio networks, and some future directions for fingerprinting in DSA systems.

II. Preliminaries

We present the basic steps for device fingerprinting, as well as frequency feature extraction. This is followed by an overview of transfer learning and related methods.

First, we establish a common notation to be used in discussing transfer learning methods, fingerprinting cognitive radios, and presenting experimental results:

R^i	A set of records from transmitter i
R^i_B	A set of records from transmitter i at bandwidth B
$\mathcal{R}$	A set of transmitters, $\{i, ..., j\}$
$r[x:y]$	Index into a record r, from point x to y
L	Features extracted from a set of records R
τ	Threshold for validation
$d(R^R, R^T)$	Distance metric between two sets of records
$D(R^T, \mathcal{R})$	Indicates a function to find a vector of distances, $[d(R^T, B^i), ..., d(R^T, B^j)]^\top, i, ...j \in \mathcal{R}$
$\mathcal{F}(r)$	Fourier transform of a single record, r
$cov(\bullet)$	The covariance of a dataset
$min(\bullet)$	The minimum of a vector
$\lvert \bullet \rvert$	The cardinality of a set

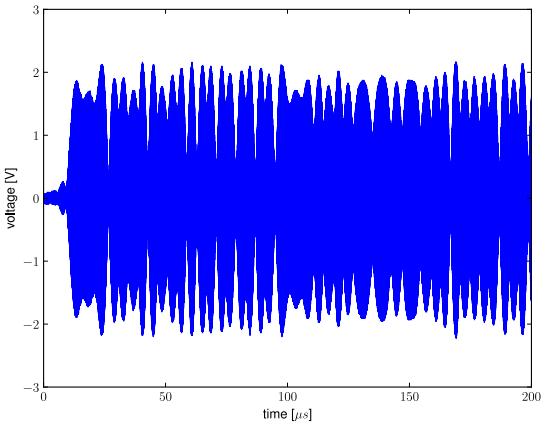

Figure 1: An example record (normalized to have zero mean and variance of one). The transient portion of the signal is on the left, and by 20 µs it has settled into steady state.

A record is a finite length discrete signal captured by an oscilloscope, see Figure 1.

A. Fingerprinting basics

When a device transmits the network must verify its identity to ensure it is following the spectrum access requirements. One or more records R^T are captured from the transmitted signal. These records must contain a portion of the signal which is constant across transmissions, generally a device identifier such as MAC addresses, or the initial part of a frame. The captured records R^T have features extracted from this constant portion, L^T, forming a test set. A database contains features extracted from records enrolled by each device to be identified. A reference set of features L^R from the asserted identity are found in the database. To determine if the device's asserted identity is valid the distance between test and reference set is calculated. If it falls below a predetermined threshold the device is accepted. In summary,

1) Extract features from captured records, $R^T \rightarrow L^T$
2) Calculate $d(L^R, L^T)$
3) Accept identity if $d(L^R, L^T) < \tau$

Mahalanobis distance is frequently used as the distance metric, but other measurements are possible. The threshold τ can be assigned per device, or to a group of devices (i.e. devices of same model, or with same transmitter parameters). The threshold is chosen based on desired false accept rate (FAR) and false reject rate (FRR) [5, Chapter 5]. The FRR is found as the percentage of time a valid user is incorrectly rejected for a given τ. Likewise, for a given τ the FAR is the percentage of time an impostor is incorrectly identified as the reference device. Adjustments to τ allow some variation between records (the performance of an identification system is very sensitive to variation in the signal, as mentioned in I). The threshold can be dynamically updated, but the basic fingerprinting method is still severely limited.

The equal error rate (EER), found by choosing τ so that $FAR = FRR$, indicates overall performance. Under some conditions the equal error rate is unsuitable for directly

comparing classifiers [5, Chapter 5], but it is an easy way to summarize performance of multiple systems.

B. Frequency feature extraction

Many possible features exist for PLI (discussed in Section VI). In this work we use frequency based features, specifically the amplitude of the instantaneous frequency of the steady state portion of the record. Features are extracted from a subset of the Fourier transform of each record. This allows reducing the dimensionality of each record, and easily applying machine learning and transfer learning techniques. Additionally, principal component analysis (PCA) is applied to the features to further reduce dimensionality.

Before extracting features from a record it is necessary to decide on a bin width (or frequency resolution), Q, found by $Q = \frac{\beta}{N}$, where β is the bandwidth covered by the features and N is the desired number of features. The record length required, X, can be found as $X = \frac{F_s}{Q}$, where F_s is the sampling frequency. For each record the start of signal, s, is found using a power threshold and refined by a matched filter with a known signal. The matched filter will produce the largest output when the start of the known signal aligns with the start of the record. The known signal can be chosen as an arbitrary record from a device. The start of signal and transient is discarded, so that only the steady state portion is used for identification. If the transient ends within d samples of the start of a record r, features can be found using $\mathcal{F}(r[o+d:o+d+X])$. The logarithm of the magnitude of the N bins centered around the carrier frequency bin form the final features.

Variation in the main lobe, as can be seen in Figure 2a, provides the best features for fingerprinting. Attenuation in the side lobes reduces the effectiveness of using these bins as features. The feature bandwidth, β, should be close to the bandwidth of the signal to include the most useful features. It is possible to choose bin width as a multiple of transmission bandwidth bandwidth, shown in Figure 2b, so that the bins have similar magnitude at all bandwidths. Superficially, this would appear to solve the problem of variation between bandwidths. However, sufficient variation exists that the transmitters cannot be accurately identified with direct comparison.

C. Transfer learning & related approaches

There are a variety of machine learning techniques to improve performance with dissimilar datasets, including representation learning, multi-view learning and transfer learning [6], [7]. Representation learning looks for underlying factors that determine the features. Use of these underlying factors can improve performance beyond what the raw features would provide. Typically this is done with deep learning techniques and neural networks. Multi-view learning describes a problem where there are multiple datasets describing the same phenomena; e.g. observations collected at different times or locations. It attempts to combine these multiple views to describe the data better than a single view can. Each view will contain slightly different information on about the phenomena, and by combining views a more complete picture is had. Transfer learning encompasses a variety of techniques to improve performance in several different problem types. Transfer learning attempts to transfer information gained from one dataset to a similar but unrelated set of data. An overview of techniques is given in [7]. It allows dealing with non-stationary distributions and boosting performance of the dataset of interest (the target) by incorporating information from another dataset (the source). From the perspective of a DSA system, the source dataset could be enrolled data from the device under test (DUT), and target data would be test data at a different bandwidth

One of the simplest transfer learning techniques is to re-weight samples in the dataset [7]. This is suitable when the source dataset only differs slightly from the target. In [8] older samples are assigned a lower weight in the training data, as are samples with an ambiguous classification. This allows adapting to gradual changes in signal characteristics. Generally, the target has less data available than is necessary for good performance. Several methods allow combining datasets, so that a larger source dataset with a "similar" distribution is used to improve performance. Some techniques also allow the source dataset to use an entirely different feature set from the target data.

Representation learning takes a different approach, and has found some success when applied to transfer learning problems. Typical machine learning algorithms require considerable effort to choose features. Although there may be hundreds of features available in a machine learning problem, each contributes very little information. Very high dimensional data needs to be reduced for most algorithms to handle it efficiently, and finding the features that best represent the data can be a time consuming task. In a very large dataset the distribution of features may be largely determined by a few "underlying factors" or "latent variables" [6]. Determining what these factors are is not straightforward. Representation learning tries to avoid the need for a human to choose the best features by applying deep learning and other techniques to reduce a high dimensional dataset to a small number of features.

Lastly, multi-view learning is useful when a source and target dataset come from a common source. In [9] a multi-view learning approach is used to allow location fingerprinting using outdated training data. The objective of location “fingerprinting” is to determine a user’s location based on received signal strength (RSS) values. Given a set of training data consisting of location coordinates and corresponding received signal strength (RSS) values at a previous time and the current time a manifold co-regularization problem is solved to find a transformation between the datasets. This allows using the complete training data from the past time while requiring data from only a few training points at the current time.

All three of these approaches are closely related. The technique outlined in this paper is based on multi-view learning, but can be seen as transfer learning where the source is the enrolled data, and the target is a set of records captured at a different bandwidth. Using current fingerprinting methods the target data cannot be identified with a high degree of accuracy. With a transfer learning approach, it should be possible to identify records at a target bandwidth using features from records at a different bandwidth.

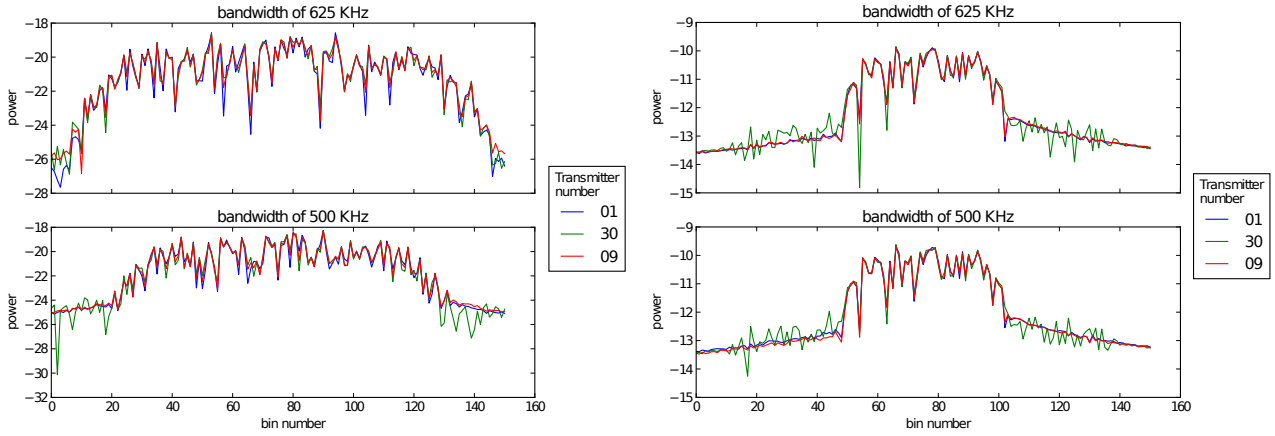

(a) (b)

Figure 2: Average of features for three transmitters when broadcasting at two different bandwidths. (a) Features using a constant bin width, note the difference between bandwidths. (b) Bin width calculated as a multiple of the bandwidth. Note the large variation in transmitter 30: although the features are more similar than in (a), identification is not possible with a direct comparison.

III. System & threat models

A dynamic spectrum access (DSA) system using cognitive radios was described in Section I-A. We now describe how fingerprinting may be applied to such a system, and how an adversary would try to subvert it.

A. System model

In the DSA network every user is allocated spectrum by a central authority. Users without authorization are blocked from using the network by jamming their signal. Each user is assigned spectrum with certain restrictions including time, duration, and transmit power. Users that violate these rules receive less bandwidth, or have their transmission jammed. This ensures selfish users will not attempt to use more than their fair share of the spectrum.

An identifier unique to each device is used to track spectrum access. This must form a recognizable part of each transmission, so that the network can determine which transmissions are valid.

B. Adversary

An adversary may violate the rules on spectrum access, or attempt to use the network without authorization. In either case this behavior will degrade performance of the network for legitimate users. In the case of a per user identifier, such as MAC addresses, it is quite easy for an attacker to impersonate a legitimate device by copying the identifier. An adversary can also record the portion of the signal used for identification, and replay it when more complex identifiers are used.

As described in Section II-A fingerprinting allows identifying devices. An attacker's transmission will have characteristics different from the device whose identity is taken. Unfortunately, current techniques are limited to comparing transmissions at the same bandwidth. Knowing this, an adversary could change bandwidth to one where the stolen identifier has not been used, and the system would be unable to refute the attackers identity. In the next section we propose a robust fingerprinting technique to fix this by allowing identification of a device at a bandwidth different from the one it has enrolled at.

IV. Robust fingerprinting

Physical layer identification is possible due to variation in the hardware of each device. Age, model, and manufacture all cause small differences in the output signal. As SDRs changes transmission parameters in software, it is to be expected that the signal variation due to hardware does not change significantly. In this case, the relationship between transmitters should be fairly consistent across bandwidths. Loosely speaking, if $d(R_1^i, R_1^j) >> 0$ then $d(R_2^i, R_2^j) >> 0$. Similarly if $d(R_1^i, R_1^j) \approx 0$ then $d(R_2^i, R_2^j) \approx 0$

Using this observation, we present the steps to extract a second set of features which are mostly invariant to changes in bandwidth. This is followed by determining the number of devices that must be enrolled at each bandwidth, and a brief description of another transfer learning approach for comparison.

A. Subspace feature extraction

Features at each bandwidth are in a subspace of the larger space of features for all possible bandwidths, as depicted in Figure 3. The bandwidth dominates the features, while the variation in each transmitter contributes relatively little. For this reason, relations between transmitters each bandwidth will be consistent, although useful comparisons between transmitters at different bandwidths are not possible. We propose a method based on a set of fixed transmitters which are used to describe each record (or set of records) within the subspace for data at a bandwidth. Features derived this way will describe a device at different bandwidths in a nearly-constant manner.

When a device asserts an identity at a bandwidth which has no enrolled data for the device its identity cannot be verified with standard fingerprinting techniques. We call this bandwidth the target bandwidth, and choose a fixed set of transmitters $\mathcal{R}$ with data enrolled at this bandwidth. The distance of the

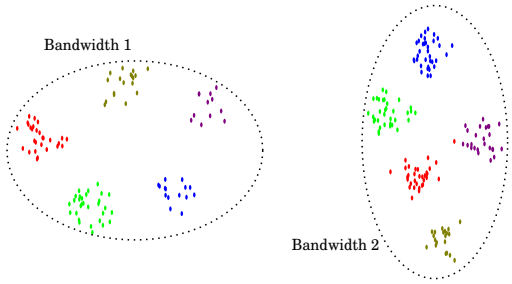

Figure 3: Example of two different subspaces in the overall feature space. The relationship between transmitters is consistent at each bandwidth, but transmitters are not directly comparable between the two bandwidths.

device under test to each device in the fixed set is found at the target bandwidth, denoted $D^T = D(L^T, L_T^{\mathcal{R}})$. At the source bandwidth the distance is measured to the same devices, and the asserted identity of the device under test: $D^S = D(L^R, L^{\mathcal{R}})$. These distances are more robust to changes, and can be compared directly for identification, $d(D^S, D^T) < \tau$.

Euclidean distance works well to find D^S and D^T. Figure 6 shows distances to nine transmitters at two bandwidths using euclidean distance. These distances vary in magnitude at the source and target bandwidths: consequently cosine distance works well for comparing D^S and D^T as it ignores magnitude. The distances to devices in $\mathcal{R}$ form features that are robust to changes in bandwidth.

In summary, given a target bandwidth with no reference data enrolled for the device under test, it can be identified using data at a source bandwidth as follows

1) Choose a set of fixed transmitters, $\mathcal{R}$ with data at target bandwidth
2) Find distances to DUT, $D^T = D(L^T, L_T^{\mathcal{R}})$
3) Find distances to asserted identity at source bandwidth, $D^S = D(L^R, L_S^{\mathcal{R}})$
4) Use $d(D^R, D^T)$ with appropriate τ to identify device

Features are only compared directly when they are from the same bandwidth. Conveniently, this opens up the possibility of using different numbers of features at each bandwidth.

B. Choice of $\mathcal{R}$

The choice of $\mathcal{R}$ will be important for best performance and to reduce the number of device-bandwidth combinations which need to be enrolled. First, we consider the size of $\mathcal{R}$. The eigenvalues of a set of data reflect the amount of variance contained in each dimension of the data. Larger eigenvalues indicate more variation in the data, while smaller values typically correspond with noise (this is used for principal component analysis: components corresponding to the largest eigenvalues are selected). The eigenvalues of the features at a single bandwidth will describe the dimensionality of the subspace at that bandwidth. In Figure 4 the first 50 eigenvalues have been plotted for records from all transmitters, at four bandwidths. Clearly the first dimension contains the most variation, and the majority of variation in the data is contained in the first five to ten dimensions. Thus, the number of fixed

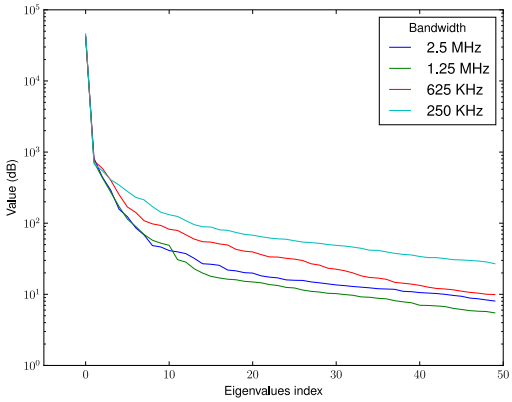

Figure 4: Sorted eigenvalues of the features at a bandwidth, using data from all devices. Most of the variance in the data is concentrated in the first dimension. The fifth eigenvalue is over two orders of magnitude smaller than the first, which suggests no more than five fixed transmitters are needed.

transmitters required will be much smaller than the number of devices enrolled in standard fingerprinting techniques.

So far we have ignored the question of which devices should be in $\mathcal{R}$. The easiest way is to randomly select devices to enroll at each bandwidth. However, randomly enrolling transmitters will allow two that are very close to each other to both be used. Since the fixed transmitters are used to describe points in a space, this could constitute redundant information.

A method to ensure the devices in $\mathcal{R}$ are not close to each other follows: let $\mathcal{R}$ denote the transmitters currently in the reference set, and i be a new transmitter. Create a new reference set $\mathcal{R}' = \mathcal{R} \cup \{i\}$. Transmitter k is removed from $\mathcal{R}'$, where k is chosen so that $min(D(L^k, L^{\mathcal{R}''})) <= d(L^l, L^j) \forall l \neq j \in \mathcal{R}''$ where $\mathcal{R}''$ is the new reference set, of the same size as $\mathcal{R}$. In other words, a transmitter is added to $\mathcal{R}$ at a given bandwidth when it is not similar to other transmitters in $\mathcal{R}$ at that bandwidth; when a device is added to $\mathcal{R}$, the device most similar to other devices in $\mathcal{R}$ is removed.

The solution is dependent on the initial choice of $\mathcal{R}$. This is a greedy algorithm and does not provide an optimum solution, nor is there necessarily a global optimum. In Section V experimental results are given for the choice of $\mathcal{R}$.

C. Another transfer learning approach

We compare our results with a transfer learning approach designed for multi-view data. In the fingerprinting problem each "view" of the data corresponds with a bandwidth.

In [10] a method of aligning multiple views is presented. Each class of data is transformed (translation and rotation) so that the mean and variance of data is the same in each view. The identification problem is based on partially unlabeled data (i.e. the true identity of the device under test is unknown), so some slight modifications are made to this method: rather than finding a per-class transform, a generic transform is found based on all datasets at each bandwidth, so that $cov(L^S) = cov(L^T)$. This amounts to a rotation of features from the source and target bandwidth, the translation is unnecessary as the data is zero mean.

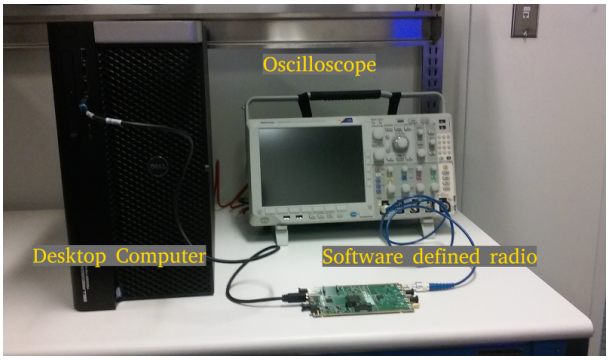

Figure 5: Experimental setup. Computer controls the SDR using GNUradio; oscilloscope captures records sent from the SDR and sends the transmission parameters to the computer.

V. Experimental evaluation

We present results of running the method on real world data. Fifty SDR devices were used with GNUradio to collect data at eleven bandwidths. Features with a bandwidth of 2 MHz were extracted from this data. Different numbers of fixed transmitters chosen randomly and as described in Section IV-B are evaluated. Although the method falls short of the ideal performance, it is better than the other methods evaluated. We present a conjecture on why the fixed transmitters accurately describe devices across bandwidths, and discuss the best way to enroll devices.

A. Transmitter parameters

The Gnuradio software provides an easy way to modulate data at different bandwidths and transmit to a SDR device. The signal was created using the QAM Mod block in GNUradio, using a quadrature amplitude modulation (QAM) signal constellation of size four modulated at 800 MHz. Eleven different bandwidths were used, with nine evenly distributed between 250 kHz and 1.25 MHz, as well as at 1.67 MHz and 2.5 MHz.

The oscilloscope used is a Tektronix DPO7354C [11], running Matlab scripts to capture data and invoke GNUradio with the correct transmit parameters for each bandwidth. The data sequence is a randomly generated sequence of bits, created beforehand and re-sent for all transmissions. The data sequence was sent repeatedly with a sufficient pause between subsequent transmissions to ensure that the oscilloscope did not trigger multiple times on the same transmission.

The data capture setup uses a desktop computer running Ubuntu 14.04 LTS. GNUradio version 3.7.10.1 was used to modulate the data for run each SDR. Fifty transmitters were used from 29 Ettus devices, shown in Table I. The N210s [12] were used with daughtercards with a single transmit frontend. The B210s [13] have two transmit frontends, and both were fingerprinted although in two cases frontend A was omitted in error, leading to fewer datasets from frontend A. This should have no impact on the data. Each transmitter frontend or daughtercard was connected to the oscilloscope with a 3 foot SMA cable(see Figure 5), and data was collected at all 11 bandwidths. A sampling rate of 20 GS/s was used, while the trigger level and oscilloscope gain was chosen for each transmitter, although the settings were reusable

Table I: Transmitters used for fingerprinting.

Model	Number of SDRs	Daughtercard or frontend	Number of datasets
N210r4	6	SBX	3
		UBX	3
B210	23	A	21
		B	23
total	29		50

for most transmitters of the same model. Once each SDR began transmitting records of length 4,000,000 were captured, consisting of noise before transmission, a transient and the steady state portion of the signal, shown in Figure 1. After discarding the transient, 3,000,000 points (150 µs) were used to extract the fingerprinting features. Records where the mean of the absolute value was much larger or smaller than normal were marked as bad, and not used in subsequent analysis. Almost all bad records occurred when the SDR first begins transmitting, and were probably due to the startup behavior of the radios. Between 481 and 521 records were collected from each transmitter at each of eleven bandwidths.

B. Feature parameters

The records captured consist of a sequence of voltage level codes. These are multiplied by the y-increment from the oscilloscope to obtain the actual voltage, then normalized so that each record has zero mean and unit variance. The features used are 150 bins with a width of 6.67 MHz. This was chosen empirically. To calculate the features 3,000,000 points are needed. These are taken beginning at 2,000 points after the detected start of the signal, to ensure that the steady-state portion of each record is used for fingerprinting.

The resulting features at each bandwidth are adjusted to have zero mean and PCA is applied to reduce the number of features to 20. This was done to speed computation times. The distance to each reference transmitter is calculated using euclidean distance, shown in Figure 6. It can be seen that the relationships between transmitters are approximately constant between bandwidths. To compare distances across bandwidths cosine distance is used.

For comparison when no transfer learning is used, the bin width was determined based on the transmission bandwidth: for a transmission at βMHz, the resulting feature width β_L was found as $\beta_L = 20e6\frac{\beta}{F_s}$. This creates the features previously shown in Figure 2b.

C. Performance of robust fingerprinting

We present results on the size of $\mathcal{R}$, and the best choice of transmitters in $\mathcal{R}$. To easily compare performance between methods, the cumulative distribution of EER is used: The threshold τ is found so that the false accept rate and false reject rate is equal. This rate is the equal error rate (EER). In a system with perfect performance it is zero, and in a practical system most should be near zero. The cumulative distribution function expresses what percentage of EERs fall below the rate on the x-axis.

Each transmitter is treated as the DUT, and all source and target bandwidth combinations are tested, excluding

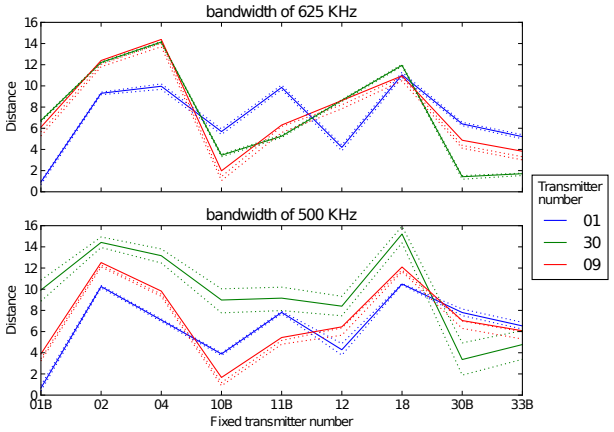

Figure 6: Distance to reference transmitters for three transmitters at two different bandwidths. The variation in features for each transmitter are largely consistent between bandwidths. Using cosine distance helps overcome the difference in magnitude. Average of distances for all records available, with dotted lines marking 10th and 90th percentile.

when the source and target bandwidth are equal. The equal error rate is found for each combination. This is repeated tentimes, randomly choosing $\mathcal{R}$ transmitters each time. The cumulative distribution of the EERs is shown in Figures 7a. Over 60% of the device, source bandwidth, and target bandwidth combinations tested have an EER below 10% with $|\mathcal{R}| = 20$. For $|\mathcal{R}| = 2$ performance is much worse: only 30% of the EERss are in this range. Of course, performance is better or worse for some target bandwidths, but this gives an impression of how the system should perform in general.

It can be seen that the results improve as the size of the fixed transmitter set increases. The improvement diminishes with increasing size, and above $|\mathcal{R}| = 10$ there is not a substantial improvement. This suggests that the relationship between bandwidths is more complex than can be described by a fixed number of points, and that ten fixed transmitters approaches the upper limit on performance.

The choice of fixed transmitters also has an impact on performance. To evaluate this, the fixed transmitters were chosen as described in Section IV-B. For each source and target bandwidth pair, the set size was increased until $|\mathcal{R}_\mathcal{T} \cup \mathcal{R}_\mathcal{S}| = M$, where M is the desired number of fixed transmitters. This gives the M 'fixed transmitters that are "most different" at both bandwidths. This was also repeated ten times with random initialization each time, and equal error rates found as before. The overall performance decreases substantially, see Figure 7b. Clearly the choice of transmitters in $\mathcal{R}$ is best done randomly. This suggests that the relation between subspaces is more complex than supposed.

D. Comparison with other approaches

Lastly we compare our method with several other approaches. The ideal performance (calculated using Mahalanobis distance with reference data is at the same bandwidth, $d(R_T^R, R_T^T)$) is shown, with over 90% of transmitters having an eer below five at any bandwidth, shown in Figure 7c. When no transfer learning method was applied, the equal error rate was found directly by calculating $d(R_S^R, R_T^T)$, using Mahalanobis distance once again. The features without any transformation applied were only tested at adjacent source and target bandwidths. Beyond this, the results were extremely poor. In a practical system, this would increase the number of bandwidths a device is required to enroll at, but even so the equal error rate is below exhibits only slightly better than random behavior, with 50% of devices having an EER below 30.

The generic transform was also tested only at adjacent bandwidths, for the same reason. Surprisingly the transfer learning method is no better than using features with bin width based on bandwidth without any transformation. This suggests the change between sub spaces alters more than the mean and covariance of the data. This is confirmed by the variation in reference distances, shown in Figure 6. Although the proposed method does not meet ideal performance, it provides a substantial improvement over existing fingerprinting techniques, and allows identifying a device enrolled at a single bandwidth at a large number of additional bandwidths.

E. Origin of similarity

Common sources of device variability include carrier oscillator offsets, variation in amplifier and filter response in the RF frontend, and, in the case of SDR, the digital-to-analog converter (DAC) used to produce the baseband signal [1]. Because we did not vary the carrier frequency and the amplifier/filter responses are likely to be small over the bandwidths considered, we conclude that the primary source of intra-device signaling variability would be attributable to an SDR's DAC.

The bandwidth of the transmitted signal is determined by the samples-per-symbol (SPS), which corresponds to the DAC outputting a given voltage for an SPS-dependent period of time. As the level of the DAC output would be primarily affected by settling time, and under the assumption that settling time is much lower than the lowest SPS duration, it is reasonable to assume that the DAC output would be invariant with respect to bandwidth for each symbol. Thus we would expect that the distance between transmitter features, which are attributable to the DAC, would also remain roughly constant across different bandwidths.

F. Enrolling devices

The number of transmitters in $\mathcal{R}$ must accurately capture the relationship between devices. More transmitters requires storing more data and enrolling devices at multiple bandwidths. Both of these actions should be minimized.

We consider a system with M devices and B possible bandwidths. Using standard fingerprinting techniques each device must have reference data enrolled at each bandwidth. This would require MB reference datasets to accurately identify each device Using the proposed method BD reference sets are required, where $D = |\mathcal{R}|$ is the number of fixed transmitters used. The amount of reference data has been greatly reduced: 550 datasets would be required for current fingerprinting techniques, In Section V-C it was shown that

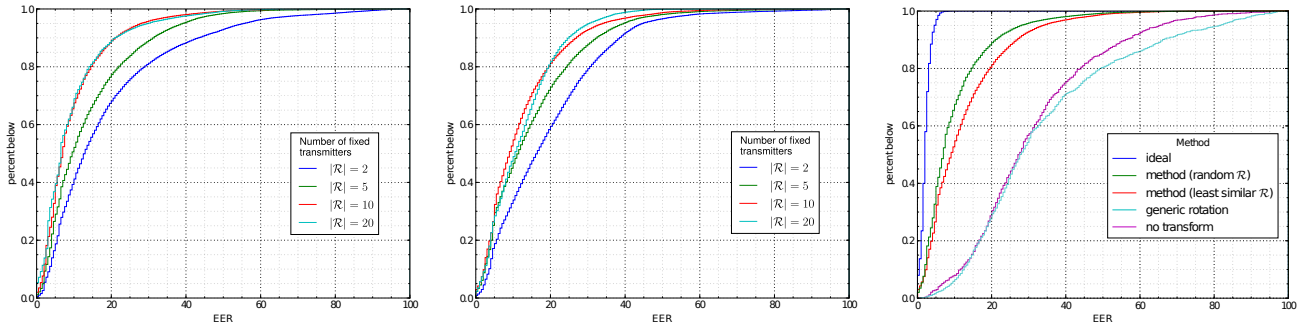

(a) (b) (c)

Figure 7: Cumulative distribution of EER of all transmitters and bandwidths, (a) EER distribution for various sizes of $\mathcal{R}$, chosen randomly. Note the diminishing returns as $|\mathcal{R}|$ increase. (b) EER distribution choosing $\mathcal{R}$ as described in Section IV-C. The performance is notably worse. (c) EER distribution comparing all methods: ideal performance; results shown in (a) and (b) with ten fixed transmitters; the generic rotation; and performance using source and target features with no transformation. The generic rotation and calculation without transformation use features with bin width based on transmitter bandwidth and only tested source and target bandwidths that are adjacent.

less than ten are required per bandwidth. As there are 50 devices in the system this is only slightly more data than current techniques require for a single bandwidth, and allows identification to occur at multiple bandwidths.

VI. Related work

We briefly cover current physical layer identification (PLI) techniques, which allow identifying devices under limited and unchanging conditions. We also cover some techniques which allow identifying devices over time, including a transfer learning approach. Channel fingerprinting is similar to PLI methods, but relies on characteristics of the channel rather than the device's hardware. Lastly, several approaches to embed unique identifiers in the signal(watermarking) are covered.

A. Physical layer identification

The state of the art for physical layer identification (PLI) is described in [1]. Features including instantaneous frequency, clock skew, transient length, timing errors, and the wavelet decomposition coefficients are used, and can be extracted from the steady state or transient portion of the record. The Fourier transform and wavelet decomposition both allow easily extracting a large number of features, and provide good performance. Results for fingerprinting a diverse set of transmitters including Bluetooth, GSM, VHF, UHF, and IEEE 802.11 transceivers are included. Most records were collected "in close proximity", using the same capture setup for each device. The systems covered have error rates (calculated as percentage of incorrectly classified records) from less than 1%, to nearly 30% with devices of the same model and manufacture.

In [4] channel equalization is used to compensate for differences between two capture setups. This decreases the EER from over 40% (nearly random) to less than two percent. Wang et al. propose a theoretical model for wireless PLI in [14]. The side lobes of the power spectrum of the signal are identified as having the most significant variation. However, in practice the attenuation in the side lobes causes significant variation between records and makes them a poor feature. It was confirmed that the fingerprint is very dependent on the channel and hardware used (e.g. increasing the distance between transmitter and receiver from 1m to 6m prevented the identification of devices until the fingerprint database was updated with records at the new distance). This matches other research [1], which has found that frequency based fingerprints are affected by changes in the channel.

B. Transfer learning for fingerprinting

In [8] a transfer learning method is applied to PLI. A clustering algorithm is used to identify users claiming multiple identities, and a transfer learning step is included to update the information about each device at each time instance. The dynamic nature of cognitive radio is not directly addressed. Although it is specifically applied to the primary user emulation (PUE) attack on a cognitive radio, the method is applicable to any fingerprinting problem where device fingerprints experience slight changes over time. The method corrects for gradual drift in features over time, not the sudden change caused by a change in bandwidth. The problem of tracking signals that vary over time has also been presented in other works, e.g. in [15] an adaptive thresholding technique was proposed.

C. Other physical layer approaches

In addition to device fingerprinting other physical layer characteristics can be identified. These approaches are based on the characteristics of the wireless channel and physical location of the device, rather than device hardware.

In [16] Xiong and Jamieson propose using angle of arrival to detect attacks. An antenna array is used to track angle of arrival of incoming packets. An attack is identified when messages claiming to be from the same device arrive from multiple directions. The angle of arrival would be very difficult to forge, but very easy to work around if the attacker has the ability place a rogue transmitter where desired. A similar approach using RSS is explored by Yu et al. in [17]. RSS values from several transmitters in the cognitive radio network are used. While it is relatively easy to alter RSS values at a single transmitter, it is much more difficult to alter them for all transmitters in the network. At the very least it would require a directional

antenna, and approximate knowledge of the placement of all devices.Simulation shows that a network of devices report RSS values can identify an attacker, as long as it is not placed near the device it is impersonating.Although channel characteristics are difficult to forge, both methods fail if the attacking device can be placed near the device it is impersonating.

In [18] Liu et al. evaluate wireless link signatures for identifying the PU. Some kind of initial knowledge of what the PU's link signature is is required, so they propose placing a helper node near the primary user. The helper node transmits when the channel is vacant, and allows secondary user (SU) to extract the link signature. The helper node transmission encounters some problems: it must be resistant to replay attacks, placed close enough to the PU to have a similar link, and be able to detect the PU's transmissions.

D. Signal embedding

Lastly, several approaches based on watermarking signals have been proposed [19]–[21]. A watermark identifying the user is embedded in each transmission. It is important that the embedding not require any changes to the transmitted signal, so that legacy systems are not impacted. To a receiver the watermarked signal should appear the same as the signal without the watermark. At the same time another cognitive radio should be able to easily extract the licensing information.

In [19] Vireshwar et al. propose embedding authentication sequence using modifications to the frequency offset of the signal. Frequency offset is compensated for at the receiver, so embedding licensing information in this way does not interfere with the normal operation of the system.

A similar proposal is given in [20] to embed information in second-order cyclostationary features. However, to avoid the need to wait for a large amount of data, Jin et al. [22] propose to embed the hash in small modifications to the QAM constellation. This can be observed immediately by other transceivers. A proposal to use the cyclic prefix is given in [21]. Simulations show that it has a much smaller impact on the bit error rate (BER) than [22], however experimental results using a SDR show a substantially higher BER than the simulation suggested.

Embedding licensing information in the QAM constellation is evaluated in [23], as well as a method exploiting error correction coding. Many systems use Reed-Solomon coding to correct errors. In a system with high signal to noise ratio (SNR) the error correction coding may not be needed to correct channel errors, as it would be possible to modify some number of bits of the signal with no error visible at the receiver.

These methods allow identifying transmitters in a way that is compatible with legacy systems. However, they are susceptible to key theft, and require additional infrastructure to distribute, manage, and revoke keys. Although the modifications to the signal are invisible to receivers, they will have an impact on performance, including bit errors and SNR.

VII. Conclusion

DSA requires new methods to prevent selfish users and attackers. We have shown that it is possible to identify cognitive radio devices with changing parameters using a comparison with a set of fixed reference transmitters. This allows current PLI methods to be applied to DSA systems, and substantially reduces the amount of training data needed. This robust fingerprinting method allows identifying devices with no modification to legacy systems, and very little overhead. The method also allows using different types of features at each bandwidth, although the improvement this provides remains to be experimentally verified. Additional investigation is required to confirm that the proposed method can be applied to changes in carrier frequency and modulation type.

References

[1] B. Danev, D. Zanetti, and S. Capkun, "Types and origins of fingerprints," in *Digital Fingerprinting*, C. Wang, R. M. Gerdes, Y. Guan, and S. K. Kasera, Eds. Springer, 2016, ch. 1.

[2] C. Wang, R. M. Gerdes, Y. Guan, and S. K. Kasera, "Digital fingerprinting," 2016.

[3] E. Hossain, D. Niyato, and Z. Han, *Dynamic spectrum access and management in cognitive radio networks*. Cambridge Univ. Press, 2009.

[4] B. Danev, S. Capkun, R. Jayaram Masti, and T. S. Benjamin, "Towards practical identification of hf rfid devices," *ACM transactions on Information and System Security (TISSEC)*, vol. 15, no. 2, p. 7, 2012.

[5] R. Bolle, *Guide to Biometrics*. Springer, 2004.

[6] Y. Bengio, A. Courville, and P. Vincent, "Representation learning: A review and new perspectives," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.

[7] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. on knowledge and data engineering*, vol. 22, no. 10, pp. 1345–1359, 2010.

[8] Y. Sharaf-Dabbagh and W. Saad, "Transfer learning for device fingerprinting with application to cognitive radio networks," 2015. [Online]. Available: http://arxiv.org/abs/1508.06614

[9] S. J. Pan, J. T. Kwok, Q. Yang, and J. J. Pan, "Adaptive localization in a dynamic wifi environment through multi-view learning," in *AAAI*, 2007, pp. 1108–1113.

[10] M. Harel and S. Mannor, "Learning from multiple outlooks," *arXiv preprint arXiv:1005.0027*, 2010.

[11] Tektronix, Inc., "DPO7000 Series Datasheet," 2017, datasheet.

[12] Ettus Research, "USRP N210 Datasheet ," 2012, datasheet.

[13] ——, "USRP B200/B210 Specification Sheet ," 2017, datasheet.

[14] W. Wang, Z. Sun, S. Piao, B. Zhu, and K. Ren, "Wireless physical-layer identification: Modeling and validation," *IEEE Transactions on Information Forensics and Security*, vol. 11, no. 9, pp. 2091–2106, 2016.

[15] R. M. Gerdes, M. Mina, S. F. Russell, and T. E. Daniels, "Physical-layer identification of wired ethernet devices," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 4, pp. 1339–1353, 2012.

[16] J. Xiong and K. Jamieson, "Securearray: Improving wifi security with fine-grained physical-layer information," in *Proceedings of the 19th annual international conference on Mobile computing & networking*. ACM, 2013, pp. 441–452.

[17] R. Yu, Y. Zhang, Y. Liu, S. Gjessing, and M. Guizani, "Securing cognitive radio networks against primary user emulation attacks," *IEEE Network*, vol. 29, no. 4, pp. 68–74, 2015.

[18] Y. Liu, P. Ning, and H. Dai, "Authenticating primary users' signals in cognitive radio networks via integrated cryptographic and wireless link signatures," in *Security and Privacy (SP), 2010 IEEE Symposium on*. IEEE, 2010, pp. 286–301.

[19] V. Kumar, J.-M. Park, and K. Bian, "Blind transmitter authentication for spectrum security and enforcement," in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 787–798.

[20] L. Yang, Z. Zhang, B. Y. Zhao, C. Kruegel, and H. Zheng, "Enforcing dynamic spectrum access with spectrum permits," in *Proceedings of the thirteenth ACM international symposium on Mobile Ad Hoc Networking and Computing*. ACM, 2012, pp. 195–204.

[21] X. Jin, J. Sun, R. Zhang, and Y. Zhang, "Safedsa: Safeguard dynamic spectrum access against fake secondary users," in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 304–315.

[22] X. Jin, J. Sun, R. Zhang, Y. Zhang, and C. Zhang, "Specguard: Spectrum misuse detection in dynamic spectrum access systems," in *Computer Communications (INFOCOM)*. IEEE, 2015, pp. 172–180.

[23] X. Tan, K. Borle, W. Du, and B. Chen, "Cryptographic link signatures for spectrum usage authentication in cognitive radio," in *Proc. of the ACM conference on Wireless network security*. ACM, 2011, pp. 79–90.