Wireless infrastructure protection using low-cost radio frequency fingerprinting receivers

*Benjamin W. Ramsey[1], Tyler D. Stubbs[1], Barry E. Mullins[1], Michael A. Temple[1], Mark A. Buckner[2]

[1]Department of Electrical and Computer Engineering, Air Force Institute of Technology, Wright-Patterson Air Force Base, Ohio 45433, USA
[2]Radio Frequency Communications and Intelligent Systems Group, Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831, USA

*Corresponding author: Benjamin W. Ramsey (benjamin.ramsey@afit.edu)

ABSTRACT

Low-data-rate wireless networks incorporated in critical infrastructure applications can be protected through 128-bit encryption keys and address-based access control lists. However, these bit-level credentials are vulnerable to interception, extraction and spoofing using software tools available free of charge on the Internet. Recent research has demonstrated that wireless physical layer device fingerprinting can be used to defend against replay and spoofing attacks. However, radio frequency (RF) fingerprinting typically uses expensive signal collection systems; this is because fingerprinting wireless devices with low-cost receivers has been reported to have inconsistent accuracy. This paper demonstrates a robust radio frequency fingerprinting process that is consistently accurate with both high-end and low-cost receivers. Indeed, the results demonstrate that low-cost software-defined radios can be used to perform accurate radio frequency fingerprinting and to identify spoofing attacks in critical IEEE 802.15.4-based infrastructure networks such as ZigBee.

KEYWORDS

Radio Frequency Fingerprinting; Physical Layer Security; WPAN, Spoofing; ZigBee Networks

1. Introduction

Low-cost, low-data-rate wireless connectivity is pervasive in critical infrastructure applications. IEEE 802.15.4-based wireless personal area networks (WPANs) operate in one-quarter of the surveyed wireless industrial control systems [1], communicate with tens of millions of smart meters [7] and are trusted components in numerous civilian and military healthcare facilities [12,21]. Security in such systems is often an afterthought, exposing critical WPANs to malicious attacks. A recent analysis of WPANs in ten U.S. cities revealed that healthcare and utility control networks operate with faulty security or none at all [17]. The threats to the critical infrastructure and other WPAN applications

[15,20] are ever increasing as open source attack tools such as KillerBee [24] and Api-do [9] become more sophisticated.

WPAN security is challenging due to the cost, power and computational constraints levied on IEEE 802.15.4-based hardware. Secure, albeit computationally-intensive, intrusion detection algorithms have been developed for high-power networks, but they are impractical for WPAN applications. While network-layer encryption is a viable option for critical networks, attackers can readily extract keys from inexpensive WPAN hardware when tamper resistance is not a design priority [8].

A promising solution for securing WPANs without placing additional burden on end devices is radio frequency (RF) fingerprinting. In such a system, an "air monitor" passively observes WPAN packets and identifies message spoofing (e.g., packet replay attacks) through device-unique radio frequency fingerprints. Wireless device classification accuracy exceeding 99% has been demonstrated using high-end signal collection receivers (with per unit cost exceeding USD 50,000) that include a 4 Gsps oscilloscope [3], 8 Gsps oscilloscope [4], 50 Gsps oscilloscope [16], a 95 Gsps Agilent E3238S signal intercept system [5,19], and an Agilent PSA E4448A spectrum analyzer combined with a 4Gsps oscilloscope [22]. The high cost of these signal receivers prohibits their use in practical radio frequency fingerprinting systems. Thus, techniques developed using high-end receivers must be successfully transitioned to low-cost (less than USD 2,000) hardware such as universal software radio peripheral (USRP) receivers. Transient-based fingerprinting requires at least 4 Gsps [3,4], which is not possible with USRP receivers that have a 25 Msps limit. However, spectral fingerprinting using wireless preambles has recently been demonstrated with USRP receivers [22,23]. Initial results suggest lower device differentiation accuracy and higher receiver-specific variability with USRP receivers than with high-end receivers.

Inexpensive analog components in low-end receivers introduce noise and variability during signal reception and confound the radio frequency fingerprinting process. While some distortion is unavoidable, the underlying hypothesis in this paper is that the variability in collection center frequency and environmental noise can be mitigated through post-collection signal processing. This paper demonstrates signal processing techniques that mitigate the radio frequency fingerprinting limitations of low-cost receivers. A key experiment described in this paper employed two radio frequency receivers (a high-end National Instruments (NI) PXIe-1085 system and a low-cost NI USRP-2921 system) under identical signal collection conditions to simultaneously collect device emissions. The results demonstrate accurate device spoofing identification in scenarios involving real-world attack hardware and smart meters.

2. Radio frequency fingerprinting

The earliest radio frequency fingerprinting systems were developed by militaries to differentiate between friendly and hostile radar transmissions [11]. The costs associated with radio frequency fingerprinting have declined over the last past decades to such a degree that commercial cell phone companies often use radio frequency fingerprinting to detect device cloning [13]. In order to be commercially viable, the radio frequency fingerprinting of low-cost WPANs in critical infrastructure applications must be practical and must leverage inexpensive, small-form-factor receiver technologies.

Ur Rehman et al. [22,23] were among the first to attempt robust radio frequency fingerprinting using low-cost USRP receivers. Their fingerprints consist solely of power spectral density (PSD) features of IEEE 802.11a (5 GHz WiFi) preambles. IEEE 802.15.4 based WPANs (e.g., ZigBee) also feature a preamble at the start of every burst transmission that is amenable to radio frequency fingerprinting. However, recent work with high-end receivers [19] reports that radio frequency fingerprints based solely on power spectral density features underperform those based on time-domain features. The fundamental hypothesis is this paper is that radio frequency fingerprinting performance using USRP receivers can match the performance of high-end receivers with proper feature selection and robust processing.

Instead of using power spectral density features, a series of instantaneous time-domain features is used to enhance the relative fingerprinting accuracy of USRP receivers. The robust radio frequency fingerprinting methodology is presented in Section 4.

3. WPAN threat scenarios

The open source `zbassocflood` tool included with KillerBee [24] enables attackers to generate fake network address requests from numerous spoofed WPAN source MAC (medium access control) addresses. These requests consume the finite pool of network addresses of higher-layer WPAN protocols such as ZigBee, causing denial-of-service attacks against legitimate devices that request access.

The standard bit-level defense against such attacks is to distribute MAC address filtering throughout a network. Access lists require time-consuming administrator management and increased memory usage on already-limited WPAN hardware. In any case, these measures are ineffective if the spoofed MAC addresses are the same as those belonging to authorized devices. A list of valid in-use MAC addresses on a network is readily found by recording nearby traffic using the `zbdump` tool in KillerBee or by evesdropping with a Microchip ZENA wireless adapter. Route poisoning, fake leave requests and other disruptive packets can be made to appear from valid source MAC addresses. Fortunately, the radio frequency fingerprints of attacker transmissions do not closely match those of legitimate devices. Thus, they would be identified as being fake by air monitors – routers would then reject the traffic based on air monitor feedback and warnings would be sent to system administrators that an attack is underway.

Replay attacks can also be used to disrupt WPANs. The `zbreplay` tool in KillerBee makes it relatively easy to conduct replay attacks. For example, an attacker could observe a WPAN-based security or control system for activity that could be replicated later (e.g., to unlock a door or open a utility valve). In such an attack, WPAN traffic that initiates the action of interest is recorded and replayed at will. The ZigBee WPAN specification does not mandate sequence number checks [19] to prevent these attacks. Even if the 8-bit ZigBee sequence field is verified by the targeted network, successful replay attacks are possible after every 255 valid frames. As with the spoofing attacks described above, the radio frequency fingerprint of the message replays would reveal that they originated from an unauthorized device. The replayed messages would be rejected and warnings would be sent to system administrators.

WPAN device localization for network auditing and cyber situational awareness are active areas of research. Inexpensive open source tools [14,18] may be used by system administrators to track down attacker hardware.

4. Radio frequency fingerprinting methodology

Since the USRP sampling rate is insufficient for transient-based radio frequency fingerprinting [3,4] and recent research has demonstrated the accuracy limitation of fingerprinting based on power spectral density [19,22,23], the methodology presented in this paper leverages the instantaneous time-domain features of the wireless preamble. Robust signal processing techniques, including frequency down-conversion and baseband filtering strategies, are used to further improve performance.

4.1 Signal collection topologies

In order to compare the relative radio frequency fingerprinting performance of high-end and low-cost receivers, it was decided to control as many parameters as possible during the signal collection phase of the experiments. Table 1 lists the control parameters for the collections made using the high-end NI PXIe-1085 system and the low-cost NI USRP-2921 system. Figure 1 shows the collection topology. Six Atmel RZUSBsticks served as the fingerprinted transmitters, each transmitting 600 IEEE 802.15.4 packets toward both collection receivers at the same time. RZUSBsticks were selected as transmitters because they were the first devices to be supported by KillerBee WPAN attack tools. All previous work on WPAN radio frequency fingerprinting has investigated the CC2420 transceiver, so the selection of the RZUSBstick broadens the research to a new transceiver type (Atmel AT86RF230). During an actual spoofing attack, a malicious device would most likely be transmitting from a different location than the impersonated device and with different hardware than used in the victim WPAN. The variances in location and hardware add to the distinctiveness of an attacker's radio frequency fingerprint. The signal collection described in Table 1 and Figure 1 correspond to the worst-case scenario experienced from a radio frequency fingerprinting perspective because the transmitters differ only in subtle physical variations in their hardware due to manufacturing tolerances.

A second set of experiments conducted at Oak Ridge National Laboratory used USRP receivers to fingerprint three OpenWay CENTRON smart meters. This expands the radio frequency fingerprinting literature to a new type of WPAN hardware. The smart meter power transmissions significantly exceeded 1 mW, so the short-range line-of-sight collection configuration shown in Figure 1 was rendered impractical without significant attenuation. In order to collect smart meter transmissions without saturating the USRP receiver, a -30 dB attenuator was introduced between the collection antenna and the USRP receiver (Figure 2). The high-end collection receiver was not portable enough to move to the stationary smart meter testbed, so direct high-end versus low-cost comparisons are only conducted using data from the first collection scenario (Figure 1).

Table 1. Radio frequency collection parameters for high-end NI PXIe-1085 and low-cost NI USRP-2921 receivers.

Parameter	Value
Transmitter-receiver distance	2 m
Height above floor (Tx/Rx)	1 m
Collection time frame	Concurrent
Transmitter	Atmel RZUSBstick
Transmit power	1 mW
Transmitter orientation	Vertical USB port
Receiver antenna	3 dBi gain VERT2450
Receiver antenna orientation	Vertical

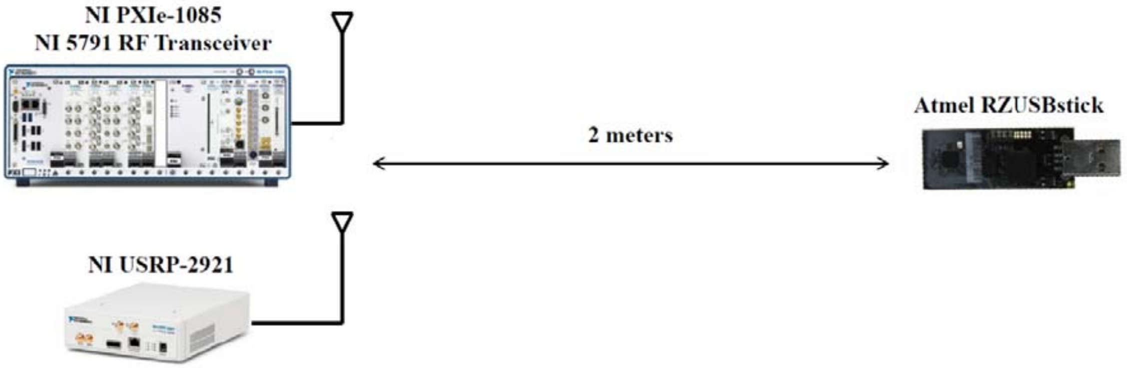

Figure 1. Collection topology for the simultaneous collection of radio frequency emissions from six Atmel RZUSBsticks using the NI PXIe-1085 and NI USRP-2921 receivers.

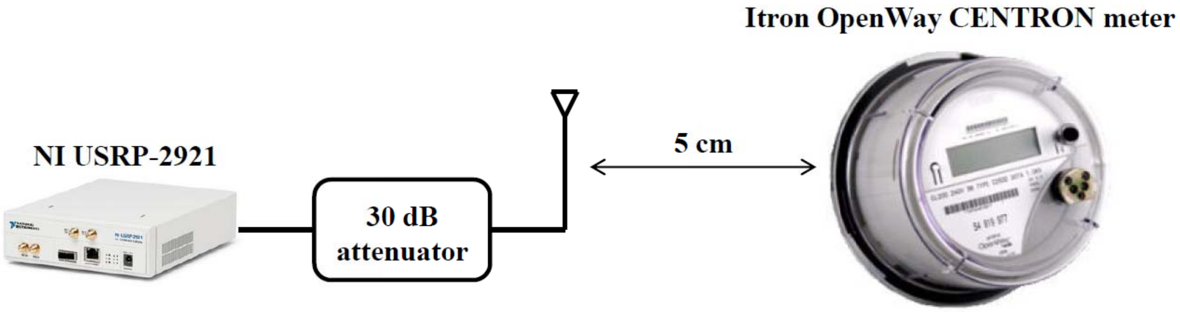

Figure 2. Collection topology for three OpenWay CENTRON Smart Meters using the NI USRP-2921 receiver.

4.2 Signal collection methodology

The signal collection methodology was consistent between the NI PXIe-1085 and NI USRP-2921 receiver systems. Both systems record radio frequency in-phase and quadrature (I/Q) data as 16-bit integers sampled at 20 Msps. This file format has the form of an interleaved array:

$$[\ I_0, Q_0, I_1, Q_1, I_2, Q_2, \ldots\ I_n, Q_n\ ]$$

where *n* is the number of collected I/Q sample pairs. The interleaved I/Q data was converted to complex values in the format:

$$[\ I_0 + iQ_0, I_1 + iQ_1, I_2 + iQ_2, \ldots\ I_n + iQ_n\ ]$$

for convenient signal processing using MATLAB. A total of 600 transmission preambles were sampled from six RZUSBsticks using the two collection receivers. Transmission detection from background noise was accomplished through amplitude-based leading edge detection using a -6 dB threshold. As outlined in the IEEE 802.15.4 standard, the first 128 µs of each transmission constitutes the preamble. At 20 Msps, the first 2,560 instantaneous I/Q samples represent the preamble region of a transmission.

Figure 3 shows a representative IEEE 802.15.4 preamble baseband response, which begins at sample number 500 and ends at sample number 3,060. The vertical dashed lines indicate the division of the preamble into 32 fingerprint regions, a process discussed in Section 4.3. The transmitter operating frequency of 2.480 GHz corresponding to IEEE 802.15.4 channel 26 was used for all collections to mitigate interference from nearby IEEE 802.11g traffic (2.401 to 2.473 GHz). The signal-to-noise ratio (SNR) was approximately 30 dB on the PXIe-1085 receiver and 24 dB on the USRP receiver.

Inter-device variability in radio frequency fingerprinting performance on USRP systems has been noted in [22]. To mitigate possible variability in collection center frequency due to clock skew, the collection center frequency was set to 3 MHz below the transmission center frequency (2.477 GHz versus 2.480 GHz). Figure 4 shows the power spectral density (PSD) of a representative transmission collected by the USRP receiver using the 3 MHz offset. The 2 MHz wide spectrum of the transmitter is notably higher than the noise floor and can be seen to be centered 3 MHz above the baseband.

The collected transmission was down-converted to the baseband via gradient-based frequency estimation performed using MATLAB. A W_{BB} = 1 MHz wide, fourth-order Butterworth filter removed background noise from outside the IEEE 802.15.4 channel, resulting in a low-noise, baseband representation of the collected transmission (Figure 5). Background noise filtering was not discussed in [22,23], which may have contributed to the erratic radio frequency fingerprinting performance that was reported.

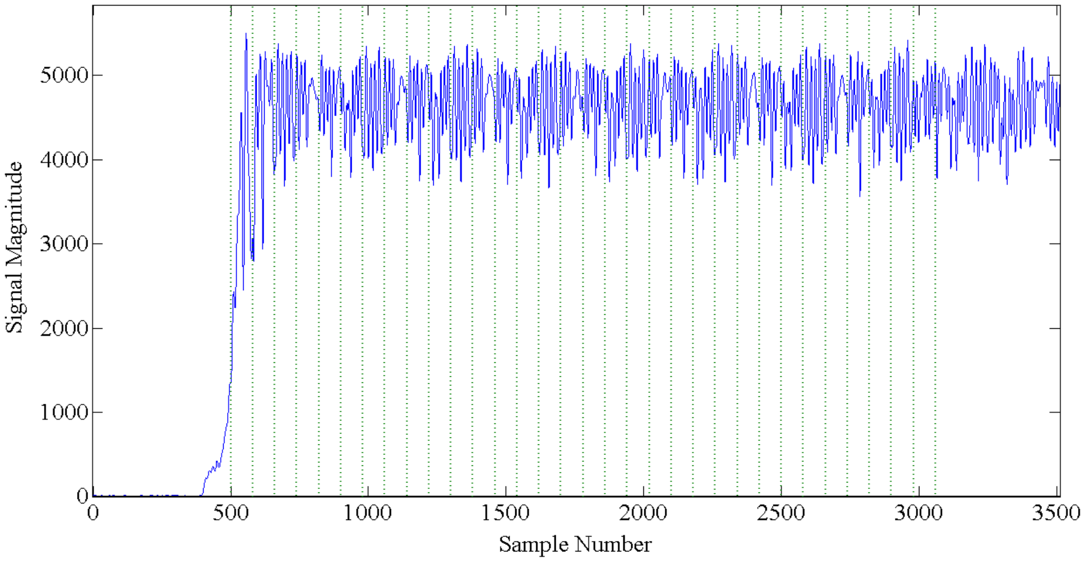

Figure 3. Provision of baseband WPAN preamble magnitude response into 32 sub-regions for radio frequency fingerprinting.

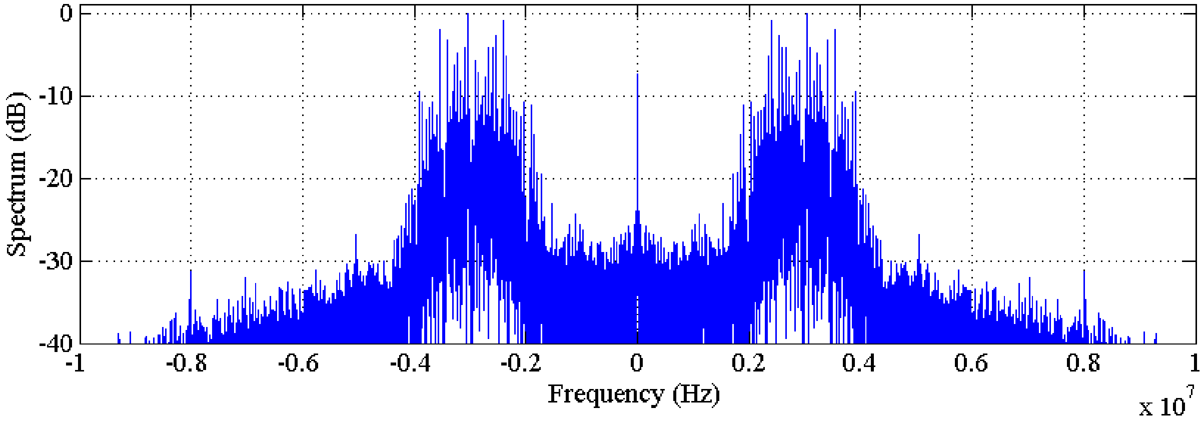

Figure 4. Normalized power spectral density of an IEEE 802.15.4 transmission collected at 20 Msps using a 3 MHz center frequency offset.

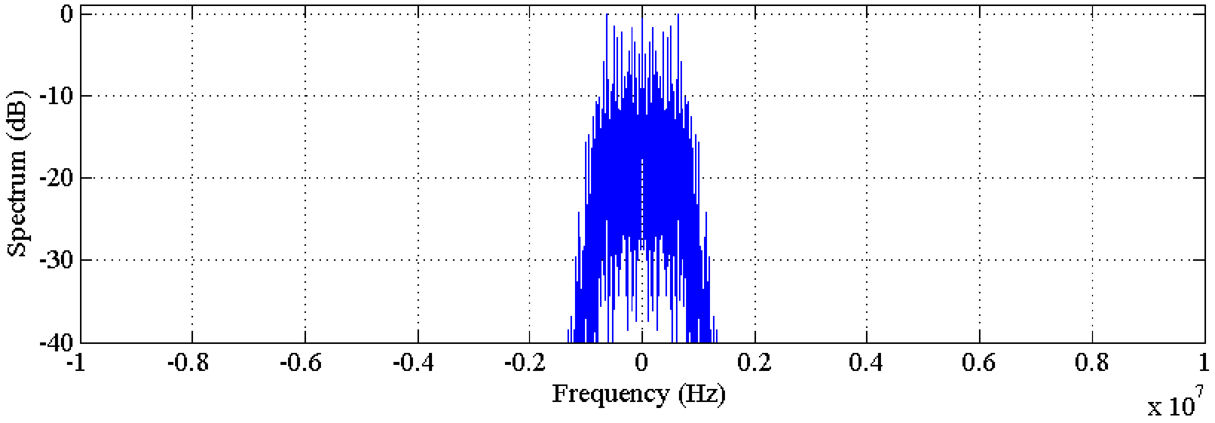

Figure 5. Normalized power spectral density of an IEEE 802.15.4 transmission down-converted to the baseband and filtered with a W_{BB} = 1 MHz wide, fourth-order Butterworth filter.

4.3 Fingerprint generation

The radio frequency fingerprint (***F***) of a signal can be derived from its instantaneous amplitude (*a*), phase (▯) and/or frequency (*f*) characteristics. More specifically, the sequences {*a*[*n*]}, {▯[*n*]}, and/or {*f*[*n*]} are generated from complex samples of the region of interest of the signal, centered (i.e., the mean is removed) and then normalized by dividing by the maximum value [19]. Instantaneous features may be computed from the I/Q characteristics of collected preambles. The instantaneous amplitude *a*[*n*] is given by:

$$a[n] = \sqrt{I[n]^2 + Q[n]^2}\ . \tag{1}$$

The instantaneous phase ▯[*n*] is given by:

$$\phi[n] = tan^{-1}\left[\frac{Q[\mathrm{n}]}{I[\mathrm{n}]}\right],\ for\ I[n] \neq 0 \tag{2}$$

and the instantaneous frequency is given by:

$$f[n] = \frac{1}{2\pi}\left[\frac{d\phi[\mathrm{n}]}{dt}\right]. \tag{3}$$

Statistical fingerprint features are generated as variance (σ^2), skewness (γ) and kurtosis (k) within specified signal sub-regions. The regional fingerprint markers are generated by: (i) dividing each sequence into N_R contiguous, equal-length sub-sequences; (ii) calculating N_S statistical metrics for each sub-sequence, plus the entire fingerprinted region as a whole (N_R + 1 total regions); and (iii) arranging the metrics in a vector of the form:

$$F_{R_i} = [\sigma^2_{R_i}\ \gamma_{R_i}\ k_{R_i}]_{1\times 3} \tag{4}$$

where i = 1, 2, …, N_R + 1. The marker vectors in Equation (4) are concatenated to form the composite vector for each characteristic, which is given by:

$$F^C = [F_{R1} \vdots F_{R2} \vdots F_{R3} \ldots F_{R(NR+1)}]_{1\times NS(NR+1)} . \qquad (5)$$

If only one signal characteristic is used (a, ⍰ or f), the expression in Equation (2) represents the final classification fingerprint. When all three signal characteristics are used, the final radio frequency fingerprint is generated by concatenating the vectors from Equation (5) according to:

$$F = \left[F^a \vdots F^\phi \vdots F^f \right]_{1\times NS(NR+1)\times NC} . \qquad (6)$$

Consistent with the radio frequency fingerprinting process introduced in [19], 32 preamble sub-regions or four regions for each of the eight repeated WPAN preamble sub-responses were employed (Figure 3). The preamble as a whole served as the 33rd region.

4.4 MDA/ML device classification methodology

Statistical radio frequency fingerprints were generated using Equation (6) for device preamble transmissions from six Atmel RZUSBStick transmitters. The resultant fingerprints were classified using a multiple discriminant analysis/maximum likelihood (MDA/ML) process in MATLAB. Multiple discriminant analysis is a straightforward extension of the Fisher linear discriminant process when the discrimination of more than two classes (devices) is required. Multiple discriminant analysis reduces the higher-dimensional input feature space with the goal of maximizing the inter-class separation while reducing intra-class spread [10]. For the six-class problems considered, the MDA/ML process projected the multidimensional radio frequency fingerprints into a five-dimensional space. The radio frequency fingerprints were classified as being affiliated with one of six possible classes based on Bayesian decision criteria using prior probabilities, probability densities and the relevant decision-making costs [6]; the costs were assumed to be the same for all the classes considered in this paper.

The MDA/ML models were developed using a k-fold cross-validation training process with k = 5 to improve reliability. This value is consistent with literature [6], which indicates that values of k = 5 and k = 10 are suitable. The best-performing model generated during the training process was used in the testing process with a previously-unseen collection of input features from each device. Model training was performed using N_{Tng} = 300 randomly-selected collected transmissions from each RZUSBstick. Testing was performed on the remaining N_{Tst} = 300 collected transmissions that were not used in the model training process.

5. RZUSBstick device classification performance

This section evaluates the relative performance of the two collection receivers with respect to device classification using radio frequency fingerprints. RZUSBstick devices were selected as transmitters because of their use in WPAN attacks launched from KillerBee and Api-do.

MDA/ML inter-device classification results were generated for all six RZUSBsticks using the high-end and low-end receivers. The classification experiments incorporated a total of 600 independent transmissions, each beginning with the IEEE 802.15.4 preamble and fifteen Monte Carlo noise realizations per preamble at each test signal-to-noise ratio (SNR). Model development used only the first 300 preambles, while testing was independently performed using the second 300 preambles. This resulted in 300 test preambles multiplied by 15 noise realizations, which is equal to 4,500 total classification decisions per device at each test SNR. This large number of trials reduced the 95% confidence intervals to within the vertical extent of the plotted markers. For visual clarity, confidence interval bars are not presented in classification plots.

5.1 Full-dimensional fingerprints at 20 Msps

Full-dimensional radio frequency fingerprints include features based on all three signal characteristics (a, ⍰ and f), three statistical features (σ^2, γ and k) and 32+1 preamble regions, for a total radio frequency fingerprint length of $N_F = 3 \times 3 \times 33 = 297$ features. The sampling rates for both collection receivers were 20 Msps. Figure 6 presents the full-dimensional classification accuracies for the six RZUSBsticks using the PXIe-1085 collection receiver while Figure 7 presents the full-dimensional classification accuracies for the same six RZUSBsticks using the USRP-2921 collection receiver. The solid black lines with circle markers in the figures show the mean classification accuracy for collections from the six transmitters.

Classification accuracies between the two collection receivers are much more consistent than those reported in [5,22] for IEEE 802.11a devices. The low-end USRP receiver classifies all six devices with an average of 90% accuracy when SNR ≥ 15 dB. The high-end PXIe-1085 system achieves an average of 90% classification accuracy when SNR ≥ 11 dB. This high-end PXIe-1085 result closely matches the findings in [5] using the high-end Agilent E3238S receiver system, where the average classification accuracy of seven CC2420 transmitters reached 90% by SNR = 10 dB. The average device classification accuracy using the USRP receiver is 9% lower than with the high-end PXIe-1085 receiver at SNR = 12 dB, but this difference narrows to 3% for SNR = 24 dB. These full-dimensional device classification results are consistent with the intuitive assumption that device classification accuracy using low-cost USRP hardware is measurably lower that the device classification accuracy of high-end signal receiver hardware. However, the difference in classification accuracy between the low-end and high-end hardware narrows to a few percent under high SNR conditions.

Given the relatively low-cost of software-defined radios such as USRP receivers, multiple receivers can be purchased for far less than a single high-end receiver. Combining radio frequency fingerprint decisions from multiple low-cost receivers may be an effective strategy to improve device classification performance. For example, if two out of three USRP-based air monitors determine the same device classification for an incoming packet, the decision based on the two concurring receivers may be more accurate than that of the dissenting receiver.

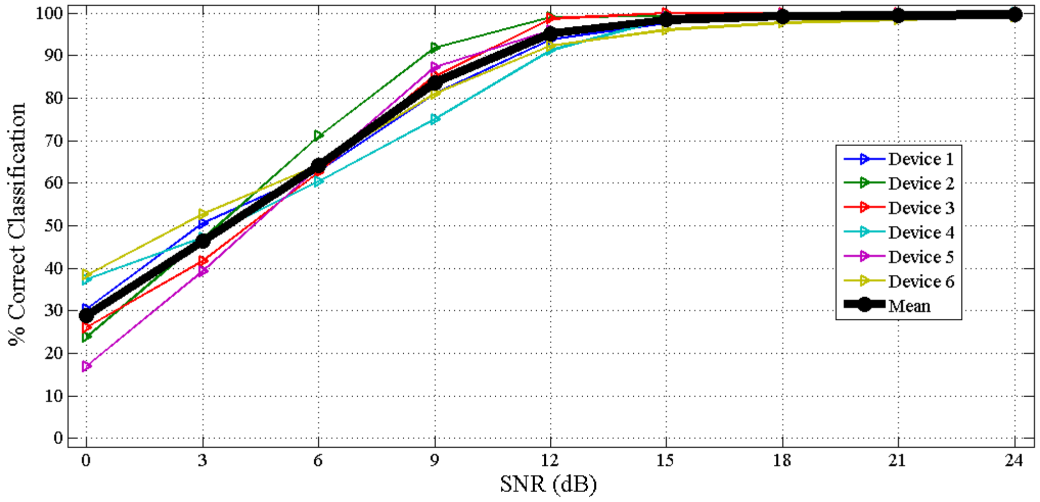

Figure 6. Classification accuracy using the NI PXIe-1085 receiver and N_F = 297 full-dimensional radio frequency fingerprints at 20 Msps.

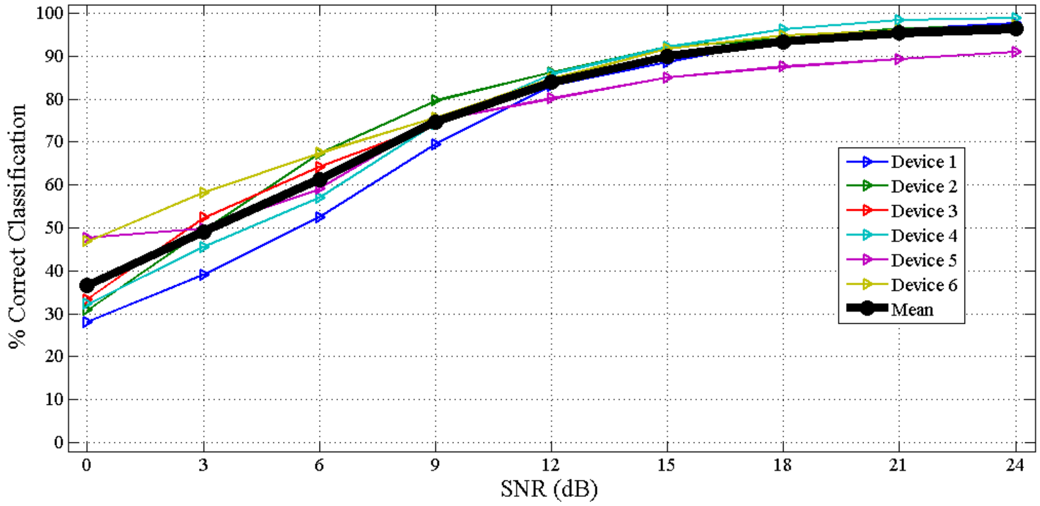

Figure 7. Classification accuracy using the NI USRP-2921 receiver and N_F = 297 full-dimensional radio frequency fingerprints at 20 Msps.

5.2 Full-dimensional fingerprints at 5 Msps

While the PXIe-1085 and USRP-2921 receivers both support sampling rates as high as 25 Msps, it is not clear that a higher sampling rate necessarily results in better radio frequency fingerprinting accuracy of WPAN transmitters. This section investigates full-dimensional radio frequency fingerprinting at the reduced sampling rate of 5 Msps.

The original radio frequency signal collections were properly decimated from 20 Msps to 5 Msps by utilizing every fourth I/Q sample and excluding the rest from the fingerprinting process. A sample rate of 5 Msps is approximately the lowest possible sample rate with which the near-baseband radio frequency collection process would work; this is because a sample rate of 5 Msps on the USRP receiver equates to a collection bandwidth spanning 2.5 MHz above and below the collection center frequency (the IEEE 802.15.4 channel width is 2 MHz). If radio frequency fingerprinting at low sample rates is effective, it decreases the hardware requirements of radio frequency air monitors deployed in operational environments.

Figures 8 and 9 present the full-dimensional classification accuracies at 5 Msps for the PXIe-1085 collection receiver and USRP-2921 collection receiver, respectively. A negligible functional difference exists between device classification accuracies at 20 MHz and 5 MHz for the high-end PXIe-1085 receiver. Similarly, the classification accuracies are functionally indistinguishable when using the USRP-2921 receiver at 20 Msps and 5 Msps. A sample rate of 5 Msps meets the Nyquist requirement for 2 MHz IEEE 802.15.4 signals and also appears to provide maximum radio frequency fingerprinting performance. Thus, low-cost radio frequency receiver hardware that supports 5 Msps may be practical for fingerprinting WPAN devices.

5.3 Phase-only fingerprints at 20 Msps

The relative importance of instantaneous amplitude, phase and frequency characteristics in device classification is reported in [19] for the high-end Agilent E3238S system and CC2420 WPAN transmitters. The results strongly suggest that instantaneous phase features are the most useful for inter-device differentiation, followed by frequency, and then instantaneous amplitudes, which tend to be the least relevant. The instantaneous phase features were robust enough that phase-only fingerprints (99 features long) were as effective as full-dimensional fingerprints with 297 features. The advantage of phase-only fingerprints is that they require calculation and processing of only one-third the number of radio frequency features. Figures 10 and 11 show the device classification accuracies of radio frequency fingerprints with only 99 instantaneous phase characteristics using the PXIe-1085 and USRP-2921 receivers, respectively.

Consistent with the high-end receiver results in [19], the phase-only radio frequency fingerprints are as effective as full-dimensional fingerprints when the high-end PXIe-1085 receiver is used as a collection receiver. Conversely, phase-only classification results with the USRP-2921 receiver significantly underperform full-dimensional radio frequency fingerprinting. The average device classification using the USRP-2921 receiver falls below 90% even when SNR = 24 dB. The results indicate that the MDA/ML device fingerprint model incorporates additional radio frequency fingerprint characteristics (more frequency or amplitude traits) when a low-end receiver is used than when a high-end receiver is used. The hypothesis is that the inexpensive analog components used in the USRP system introduce additional radio frequency fingerprint distortion that the MDA/ML model overcomes by diversifying the instantaneous characteristics given the most weight during model development.

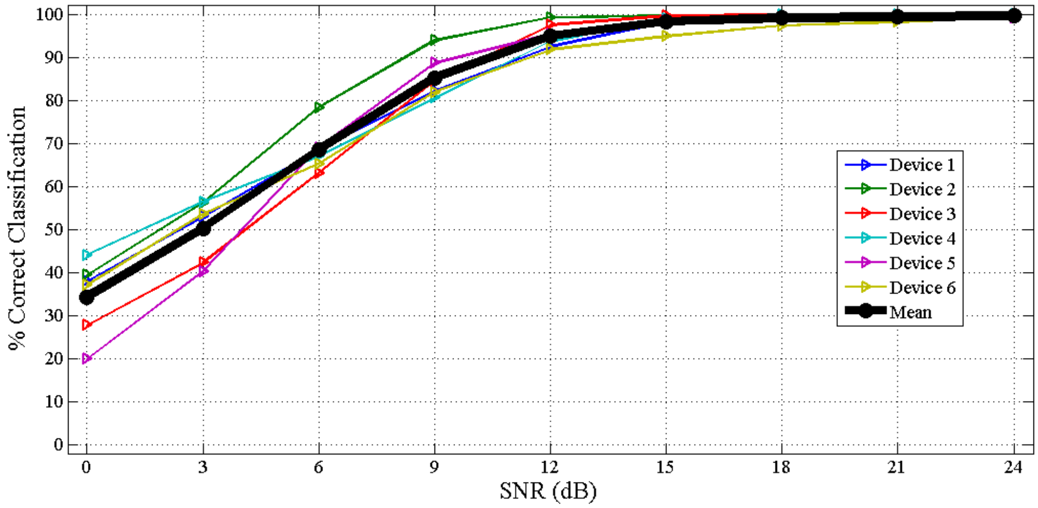

Figure 8. Classification accuracy using the NI PXIe-1085 receiver and N_F = 297 full-dimensional radio frequency fingerprints at 5 Msps.

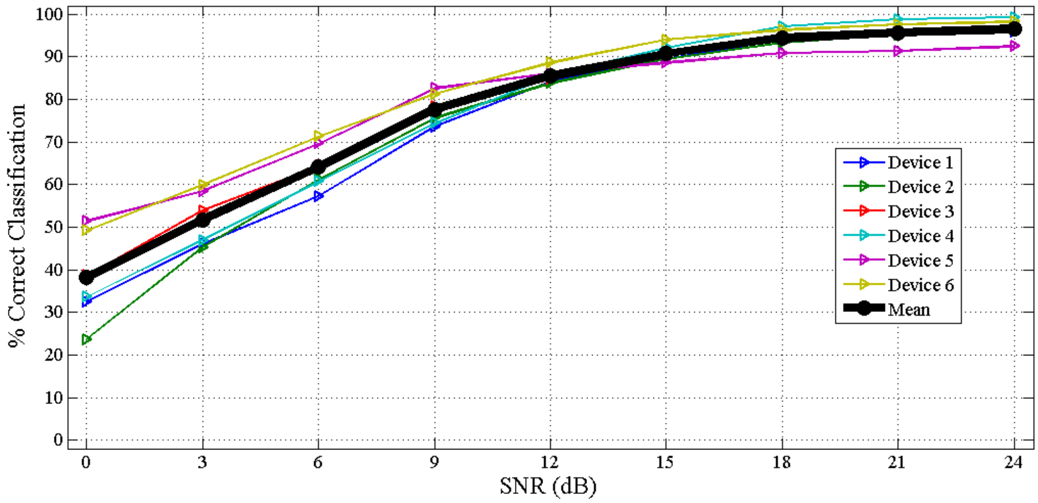

Figure 9. Classification accuracy using the NI USRP-2921 receiver and N_F = 297 full-dimensional radio frequency fingerprints at 5 Msps.

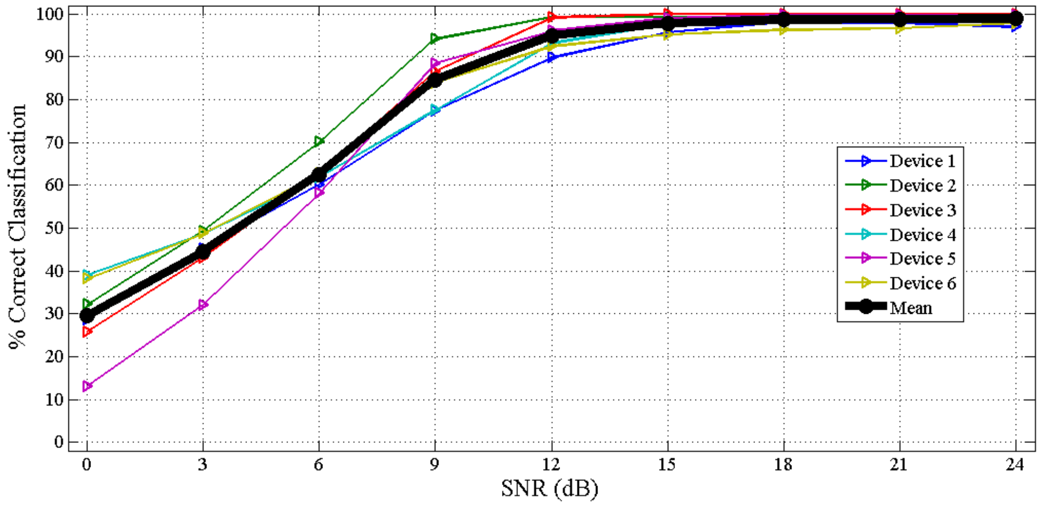

Figure 10. Classification accuracy using the NI PXIe-1085 receiver and N_F = 99 phase-only radio frequency fingerprints at 20 Msps.

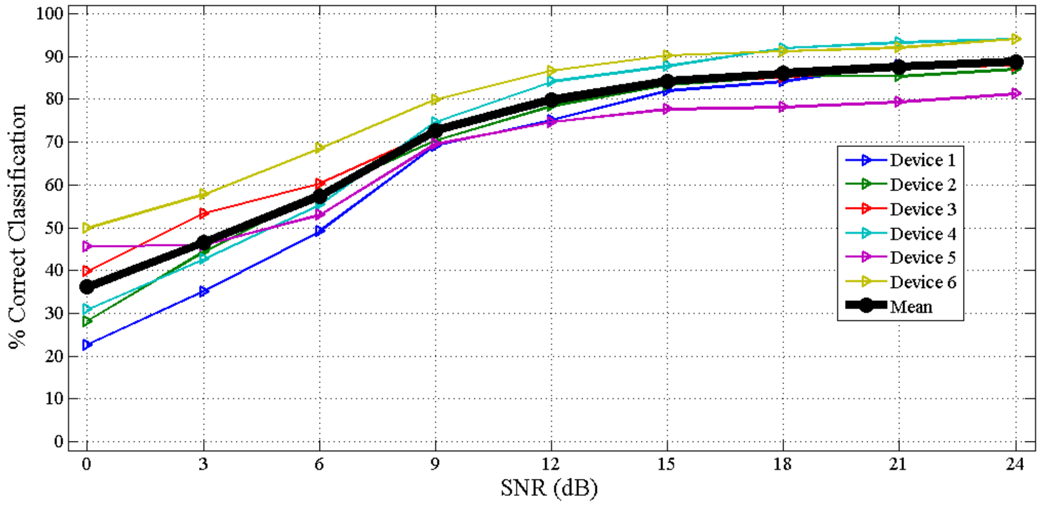

Figure 11. Classification accuracy using the NI USRP-2921 receiver and N_F = 99 phase-only radio frequency fingerprints at 20 Msps.

To test this hypothesis, device classification was performed using both collection receivers and fingerprints consisting of only one of the three radio frequency characteristics (amplitude, phase or frequency). The classification results present the relative feature relevance to MDA/ML model development on signals collected by the two receivers (Figures 12 and 13). The phase-only classification results reported in Figures 12 and 13 correspond to the classification means in Figures 10 and 11, respectively.

While the relative relevance of instantaneous radio frequency characteristics remains the same for the Agilent E3238S [19], PXIe-1085 and USRP-2921 systems (phase > frequency > amplitude), the classification accuracy achieved using any one instantaneous characteristic is significantly lower for the USRP-2921 system; the arbitrary 90% correct classification benchmark is not achieved in any case.

For example, the mean device classification accuracy using fingerprints consisting solely of phase or frequency characteristics is sufficient for near-100% accuracy for SNR > 21 with the high-end receivers. However, the device classification accuracies achieved using equivalent single-characteristic fingerprints with the USRP-2921 receiver are 10% or more lower than those achieved with high-end receivers at SNR = 24 dB. Dimensionality reduction of the fingerprints generated with low-end collection receivers is still possible through feature ranking, but trivial reduction through selection of a single fingerprint characteristic to incorporate is clearly not possible with a low-cost USRP receiver.

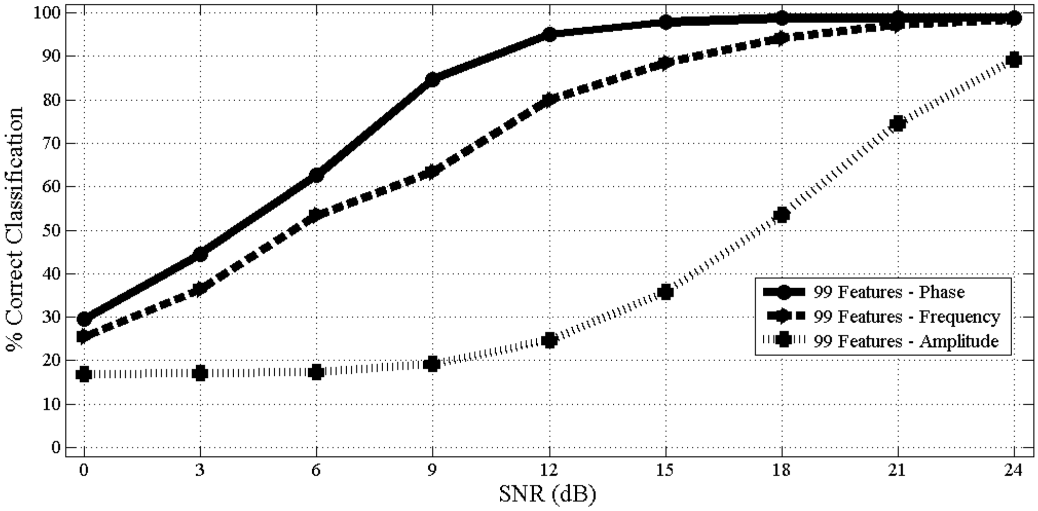

Figure 12. Classification accuracy using the NI PXIe-1085 receiver and N_F = 99 single-characteristic radio frequency fingerprints at 20 Msps.

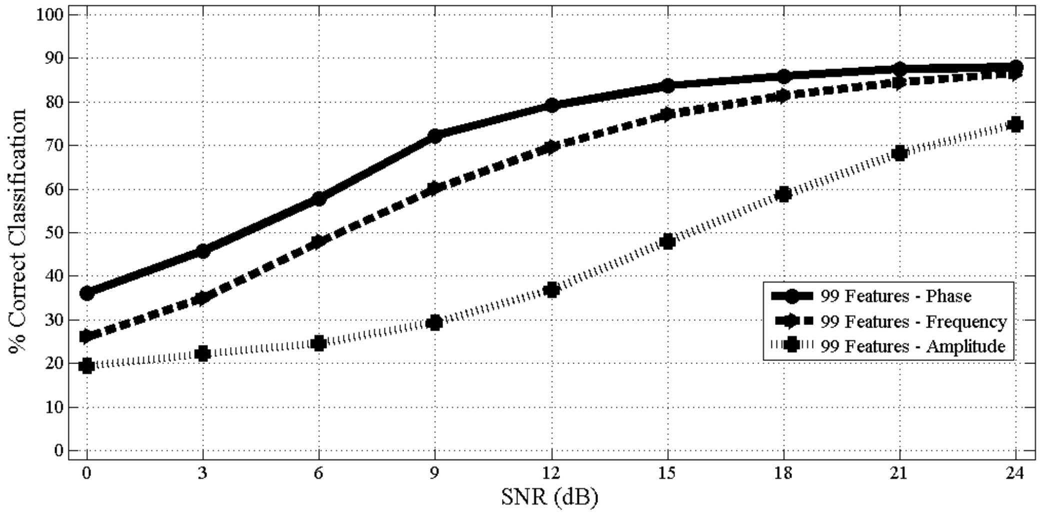

Figure 13. Classification accuracy using the NI USRP-2921 receiver and N_F = 99 single-characteristic radio frequency fingerprints at 20 Msps.

6. Smart meter classification performance

Smart meter device classification was evaluated using radio frequency fingerprints with the USRP-2921 receiver. The high-end PXI-1085 received was not portable enough to be relocated within the stationary testbed of OpenWay CENTRON Smart Meters at Oak Ridge National Laboratory. The collection topology used is shown in Figure 2. The -30 dB signal attenuator placed between the receiver antenna and USRP-2921 receiver decreased the SNR to 14 dB. The device classification results are reported for SNR ∈ [0, 12] dB because white Gaussian noise (AWGN) can only be added to the original signal collections to produce the test SNR environments, not further reduced.

Figure 14 presents the device classification results using the USRP-2921 receiver and full-dimensional fingerprints at 20 Msps. The mean device classification accuracy reaches 90% when SNR = 6 dB and increases to 96% at SNR = 12 dB. Figure 15 presents the device classification results using the USRP-2921 receiver and full- dimensional fingerprints calculated from properly decimated radio frequency collections at 5 Msps. As observed earlier for the six RZUSBsticks, the mean device classification accuracy does not fall when the sample rate decreases from 20 Msps to 5 Msps. This is additional evidence that 5 Msps is sufficient to achieve maximum radio frequency fingerprinting performance for IEEE 802.15.4 devices.

7. Device identity verification

The device classification results in the earlier sections establish that the radio frequency fingerprinting methodology is effective at inter-device differentiation. This section demonstrates the ability of the radio frequency fingerprinting methodology to detect device spoofing attacks against critical infrastructure WPANs.

As described previously, radio frequency fingerprints were generated for nine devices, six RZUSBsticks and three smart meters, using the USRP-2921 receiver for signal collection. In a typical scenario, the air monitor system would be trained using radio frequency fingerprints calculated from transmissions made by its trusted member devices. Spoofing attacks originate from untrusted devices with hardware-unique radio frequency fingerprints that do not exactly match those of trusted devices in the WPAN.

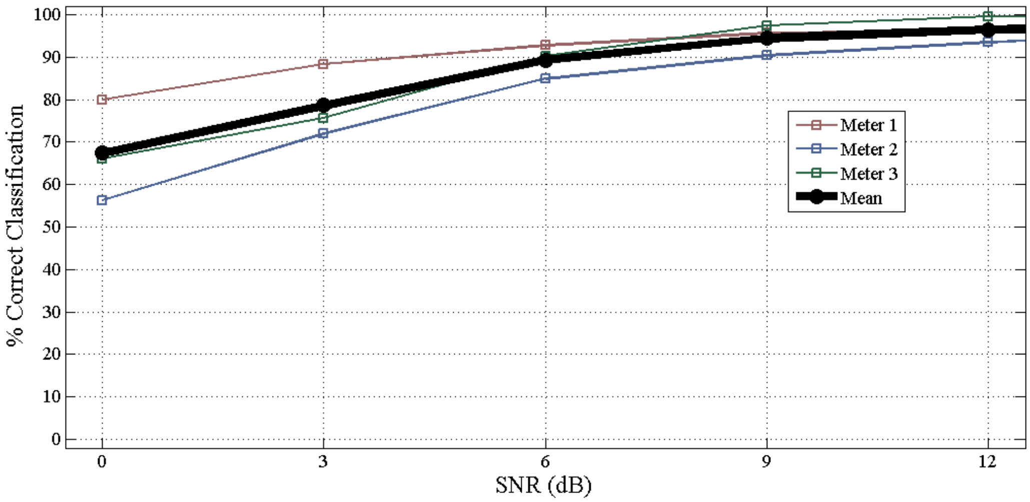

Figure 14. Smart meter classification using the NI USRP-2921 receiver and N_F = 297 full-dimensional radio frequency fingerprints at 20 Msps.

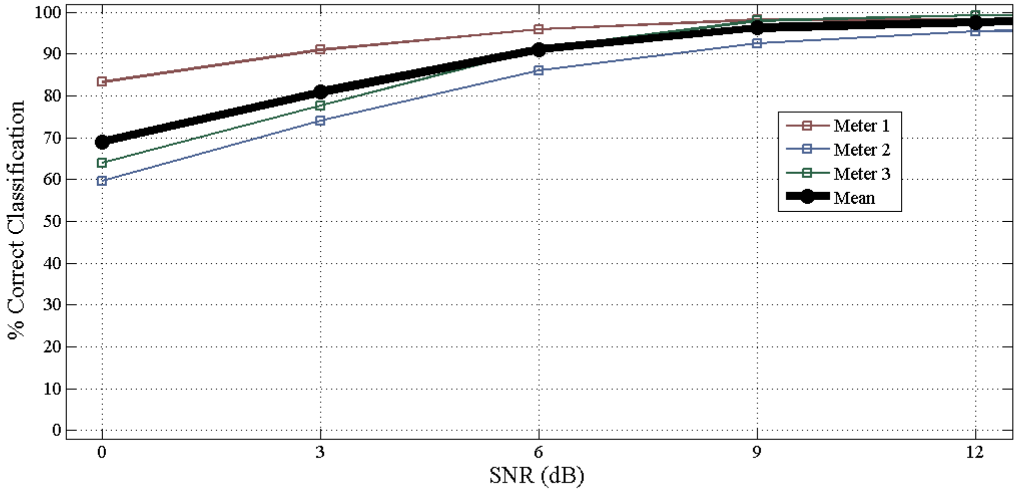

Figure 15. Smart meter classification using the NI USRP-2921 receiver and N_F = 297 full-dimensional radio frequency fingerprints at 5 Msps.

7.1 Device verification scenario

The verification methodology in [5] was adopted with a subset of authorized devices used for air monitor training and the remaining devices were used to perform spoofing attacks against each of the authorized devices. Since RZUSBsticks are a popular WPAN hardware attack platform, three RZUSBsticks were selected (Devices 1, 2 and 3 as labeled in Figures 6 through 11) to serve as the spoofing devices. The remaining three RZUSBsticks (Devices 4, 5 and 6) and three smart meters (Meters 1, 2 and 3) formed the pool of authorized WPAN devices. The combination of smart meters and RZUSBsticks is consistent with a smart grid WPAN implementation that incorporates interconnected smart meters and industrial appliances.

First, an MDA/ML device classification model was generated for the six authorized devices using N_{TNG} = 300 preamble-based full-dimensional radio frequency fingerprints each (as described in Sections 5 and 6). This created a five-dimensional Fisher projection that maximized inter-device differentiability. Radio frequency fingerprints from the three spoofing devices underwent the same Fisher projection as the authorized devices. Each of the three spoofing devices was introduced as an impersonator for each of the six authorized devices, resulting in a total of 3 x 6 = 18 spoofing scenarios. SNR was introduced for verification and spoofing rejection was assessed at SNR = 12 dB.

The posterior output variable from the MATLAB *classify* function provides the verification test statistic for a spoofing device because it impersonates each authorized device. Spoofing device verification is assessed by inputting the posterior output to the MATLAB *roc* (receiver operating characteristic) function, which yields verification performance curve data.

7.2 Device verification accuracy

Given a spoofing Device *j* (*Dj*) that presents a claimed identity of Device *I* (*Di*), two probabilities may be used to generate verification performance curves for spoofing scenarios: (i) P[*Di*|***Fi***] provides a measure of how much authorized *Di* projected fingerprints "look like" authorized *Di*; and (ii) P[*Di*|***Fj***] provides a measure of how much the spoofing device *Dj* "looks like" the authorized device *Di*. These probabilities were used to generate the results presented in Figure 16. The ROC legend for each of the 18 spoofing scenarios is in the format {*spoofing device: spoofed device*}.

The results presented in Figure 16 are interpreted as follows. The vertical axis represents the probability that the authorized device is recognized as legitimate *(True Verification Rate)* and is, therefore, accepted by the WPAN. The horizontal axis represents the probability that the spoofing device successfully impersonates the authorized device (*Rogue Accept Rate)*. A ROC curve that approaches the upper left corner of Figure 16 (TVR = 100% and RAR = 0%) indicates that a statistical threshold exists where all transmissions from an authorized device are accepted as legitimate and all spoofing attacks are rejected. Conversely, the ROC curves removed from the upper left corner indicate only imperfect spoofing detection.

In 16 out of 18 spoofing scenarios (collection of curves near the upper left corner), a threshold of TVR > 90% authorized packet acceptance results in less than RAR < 2% acceptance of spoofed packets. Spoofing detection is 100% accurate for all scenarios in which Devices 1, 2 and 3 (Dev 1, Dev 2 and Dev 3) impersonated the three smart meters. In the two most challenging scenarios ({Dev 1: Dev 4} and {Dev 3: Dev 6}), the TVR > 90% threshold results in spoofed packet acceptances of 54% and 32%, respectively. It is important to recall that the RZUSPsticks were fingerprinted under atypical challenging conditions, where many factors were controlled that would otherwise have contributed to inter-device differentiability, including device position and transmit antenna orientation. Given the challenges imposed, the robustness of this device verification process is apparent. Even WPAN devices of the same hardware type and in the same antenna orientation can be reliably verified using radio frequency fingerprints generated from signals collected by a low-cost USRP receiver.

A useful technique for visually representing RF fingerprints is through radio frequency – distinct native attributes (RF-DNA) markers, adopted here from [5]. Figure 17 shows the average RF-DNA responses for the six RZUSBstick transmitters generated using the NI USRP-2921 receiver at 20 Msps. Averages were calculated based on 400 preambles at SNR = 12 dB. Full-dimensional N_F = 297 RF fingerprints include 99 markers for each of the three statistics (variance (*var*), skewness (*skw*) and kurtosis (*kur*)) as described in Section 4.3. It is important to note that this normalized (within statistic) representation was developed to help visualize feature variation across devices. This particular normalization is not included when using radio frequency fingerprints for classification and verification.

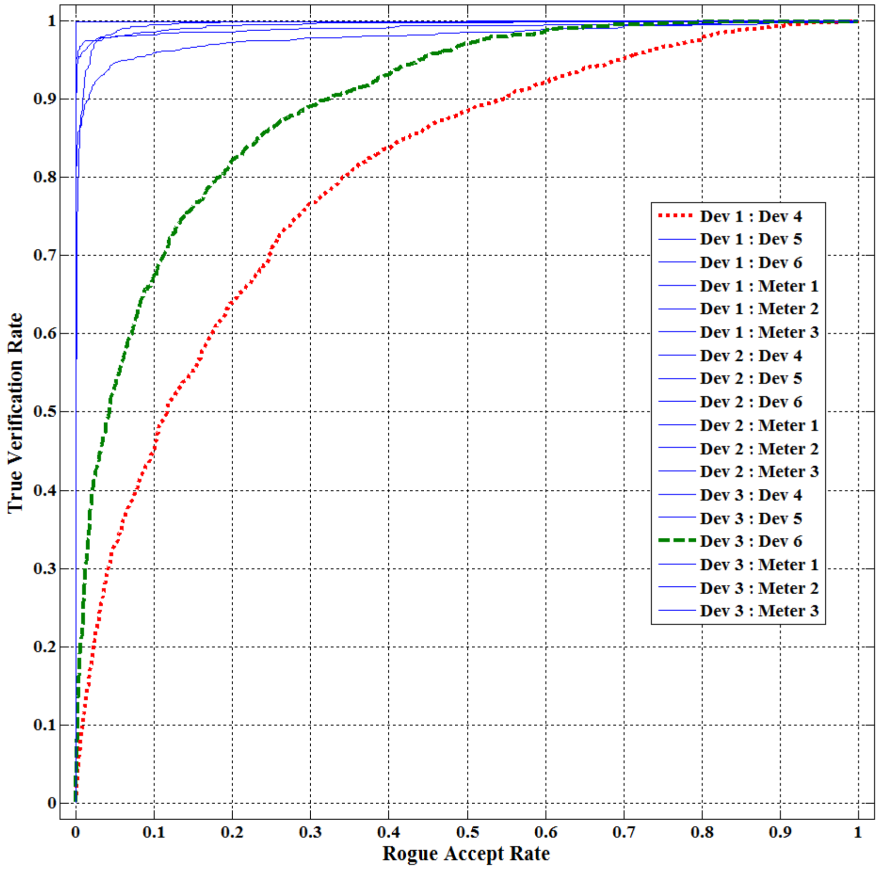

Figure 16. Device verification using the NI USRP-2921 receiver at 20 Msps and N_F = 297 full-dimensional radio frequency fingerprints at SNR = 12 dB.

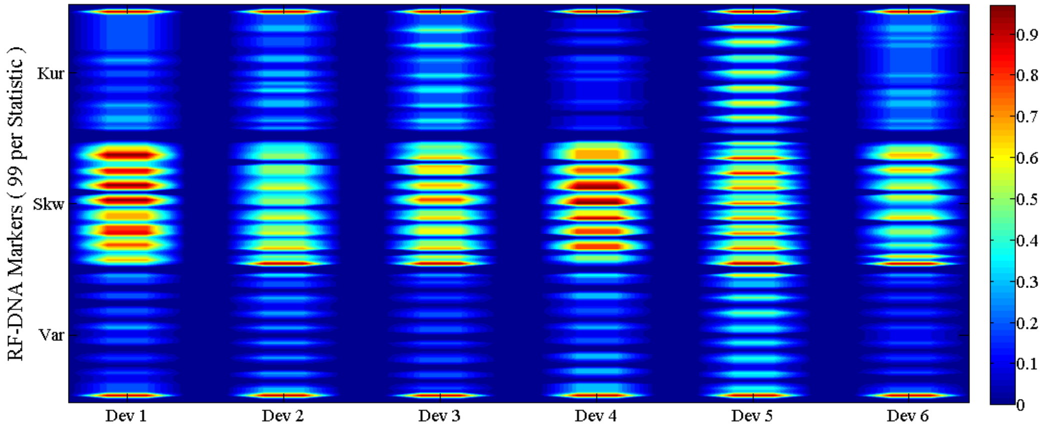

Figure 17. Average RF-DNA markers for Devices 1-6 at SNR = 12 dB.

Dev 1 and Dev 4 appear to be the most similar in the RF-DNA visualization shown in Figure 17. The similarity between these devices mirrors the spoofing detection challenge reported by the {Dev 1: Dev 4} ROC curve in Figure 16. The second most challenging spoofing scenario in Figure 16 is the impersonation of Dev 6 by Dev 3. Nevertheless, inter-device similarities are apparent from the RF-DNA markers of Dev 3 and Dev 6 shown in Figure 17.

8. Conclusions

This paper demonstrates that reliable radio frequency fingerprinting of IEEE 802.15.4 networks used in the critical infrastructure is practical using low-cost signal receivers. The distortions introduced by inexpensive analog components are mitigated by conducting signal collection with a small frequency offset and by filtering the background noise effects. The results suggest that the 25 Msps sampling rate of the NI USRP-2921 receiver is not essential to IEEE 802.15.4 fingerprinting, and that lower cost receivers supporting 5 Msps are sufficient to protect operational systems. The results also demonstrate that frequency-based fingerprint features are more relevant to device classification when a low-cost signal receiver is used as opposed to a high-end receiver. Indeed, this research represents a significant step toward realizing a practical, low-cost radio frequency fingerprinting solution.

Near-term research will continue to focus on lowering the costs associated with implementing practical radio frequency fingerprinting solutions. The Nuand bladeRF software-defined radio has a fraction of the cost of the NI USRP-2921 receiver and is a promising candidate for future work. IEEE 802.15.4 WPANs are of particular interest in future radio frequency fingerprinting experiments because of their widespread use in the critical infrastructure and the unique security challenges they pose.

Acknowledgements

The authors thank the reviewers for their insightful comments. Note that the views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, U.S. Department of Defense, U.S. Department of Energy or United States Government.

IMPORTANT NOTE TO IJCIP TYPESETTERS: I have edited the references in the paper myself. You are welcome to change the order of the references and the order of the items within a reference. However, DO NOT change any capitalization or fonts (e.g., italics) in the references below. Professor Sujeet Shenoi, Editor-in-Chief, IJCIP

References

[1] W. Boyes, Industrial wireless: All quiet on the wireless front, *Control Global*, Schaumburg, Illinois (`www.controlglobal.com/articles/2011/all-quite-on-the-wireless-front.html`), August 9, 2011.

[2] J. Cache, J. Wright and V. Liu, *Hacking Exposed Wireless: Wireless Security Secrets and Solutions*, McGraw-Hill, New York, 2010.

[3] B. Danev and S. Capkun, Transient-based identification of wireless sensor nodes, *Proceedings of the International Conference on Information Processing in Sensor Networks*, pp. 25-36, 2009.

[4] B. Danev, H. Luechken, S. Capkun and K. El Defrawy, Attacks on physical-layer identification, *Proceedings of the Third ACM Conference on Wireless Network Security*, pp. 89-98, 2010.

[5] C. Dubendorfer, B. Ramsey and M. Temple, An RF-DNA verification process for ZigBee networks, *Proceedings of the Military Communications Conference,* 2012.

[6] R. Duda, P. Hart and D. Stork, *Pattern Classification*, John Wiley and Sons, New York, 2001.

[7] Federal Energy Regulatory Commission, Assessment of Demand Response and Advanced Metering, Staff Report, Washington, DC (`www.ferc.gov/legal/staff-reports/12-20-12-demand-response.pdf`), 2012.

[8] T. Goodspeed, Extracting keys from second generation ZigBee chips, presented at *Black Hat USA* (`www.blackhat.com/presentations/bhusa-09/GOODSPEED/BHUSA09-Goodspeed-ZigbeeChips-PAPER.pdf`), 2009.

[9] T. Goodspeed, S. Bratus, R. Melgares, R. Speers and S. Smith, Api-do: Tools for exploring the wireless attack surface in smart meters, *Proceedings of the Forty-Fifth Hawaii International Conference on System Sciences*, pp. 2133-2140, 2012.

[10] T. Hastie, R. Tibshirani and J. Friedman, *The Elements of Statistical Learning; Data Mining, Inference and Prediction*, Springer, New York, 2009.

[11] N. Hu and Y. Yao, Identification of legacy radios in a cognitive radio network using a radio frequency fingerprinting based method, *Proceedings of the IEEE International Conference on Communications*, pp. 1597-1602, 2012.

[12] R. Istepanian, E. Jovanov and Y. Zhang, Guest Editorial, Introduction to the special section on M-Health: Beyond seamless mobility and global wireless healthcare connectivity, *IEEE Transactions on Information Technology in Biomedicine*, vol. 8(4)**,** pp. 405-414, 2004.

[13] D. Kapan and D. Stanhope, Waveform Collection for use in Wireless Telephone Identification, U.S. Patent 5,999,806, December 7, 1999.

[14] C. Kiraly and G. Picco, Where's the mote? Ask the MoteHunter! *Proceedings of the Thirty-Seventh IEEE Conference on Local Computer Networks Workshops*, pp. 982-990, 2012.

[15] M. Lin, J. Leu, K. Li and J. Wu, ZigBee-based Internet of Things in 3D terrains, *Computers and Electrical Engineering*, vol. 39(6), pp. 1667-1683, 2013.

[16] A. Polak, S. Dolatshahi and D. Goeckel, Identifying wireless users via transmitter imperfections, *IEEE Journal on Selected Areas in Communications*, vol. 29(7), pp. 1469-1479, 2011.

[17] B. Ramsey, B. Mullins, R. Speers and K. Batterton, Watching for weakness in wild WPANs, *Proceedings of the Military Communications Conference*, pp. 1401-1409, 2013.

[18] B. Ramsey, B. Mullins and E. White, Improved tools for indoor ZigBee warwalking, *Proceedings of the Thirty-Seventh IEEE Conference on Local Computer Networks Workshops*, pp. 921-924, 2012.

[19] B. Ramsey, M. Temple and B. Mullins, PHY foundation for multi-factor ZigBee node authentication, *Proceedings of the IEEE Global Communications Conference*, pp. 795-800, 2012.

[20] R. Roman, C. Alcaraz, J. Lopez and N. Sklavos, Key management systems for sensor networks in the context of the Internet of Things, *Computers and Electrical Engineering*, vol. 37(2), pp. 147-159, 2011.

[21] C. Swedberg, Air Force hospital eliminates equipment loss, reduces labor hours, *RFID Journal* Hauppauge, New York (www.rfidjournal.com/articles/view?8445), May 23, 2011.

[22] S. Ur Rehman, K. Sowerby and C. Coghill, Analysis of receiver front end on the performance of RF fingerprinting, *Proceedings of the Twenty-Third IEEE International Symposium on Personal, Indoor and Mobile Radio Communications*, pp. 2494-2499, 2012.

[23] S. Ur Rehman, K. Sowerby and C. Coghil, Analysis of impersonation attacks on systems using RF fingerprints and low-end receivers, *Journal of Computer and Systems Sciences*, vol. 80(3), pp. 591-601, 2014.

[24] J. Wright, KillerBee: Framework and Tools for Exploiting ZigBee and IEEE 802.15.4 Networks. Version 1.0 (code.google.com/p/killerbee), 2010.