Interference Testing on Radio Frequency Polarization Fingerprinting

Page Heller
Endpoint Security Inc.
College Station, Texas, USA
email: heller@endpointsecurityinc.com

***Abstract*— As nation state actors become more active in cyber-attacks on infrastructure, they also become more sophisticated, choosing to target product quality over a plant shutdown, thus making it harder to detect an intrusion. To make cyber-attacks harder to initiate, naturally-occurring polarization in radio frequency signals is being explored as a means of authentication that doesn't require digital data. By means of polarization mode dispersion, it is possible to protect the wireless channel (the path from sensor to access point) by identifying hostile actors who attempt to imitate authenticated devices to gain entry into a wireless network. In this article, recent test results are examined for their impact on the resilience of this type of wireless security. Specifically, performance in the case of low received-signal-strength is analyzed. Also, troublesome presence of electrical interference from a microwave oven and fans is studied.**

Keywords-cybersecurity; authentication; wireless intrusion detection; radio frequency fingerprinting; interference.

I. INTRODUCTION

As nation state actors become more active in cyber-attacks on infrastructure, they also become more sophisticated [1]. Rather than shutting down a plant, for instance, they might target product quality, which is harder to detect. Rather than using dictionary attacks, they might employ man-in-the-middle attacks or implement rogue access points because they are also harder to detect [2]. To stay out in front of the attackers requires a proactive approach that includes new systems of security for wireless edge devices.

A new form of Wireless Intrusion Detection System (WIDS) has been developed that detects wireless intruders by the signal they send, rather than simply relying upon the data that the signal contains [3] [4]. By using polarization mode dispersion, it is possible to protect a wireless channel (that is, the path from sensor to access point) by identifying hostile actors who attempt to imitate authenticated devices to gain entry into a wireless network.

In this article, recent test results are examined for their impact on the resilience of this type of wireless security. Specifically, performance in the case of low received signal strength is analyzed. Also, presence of electrical interference from a microwave oven and from fans is studied.

The remainder of this paper is organized as follows. Section II describes the test setup, including a description of a prototype device under test, and further describes a set of fixed transmitting devices emulating industrial sensors, the general environment in which tests are conducted and the condition under test, which may involve introducing a tertiary device. Section III describes the test performed with signals received that are characterized by low relative signal strength. Section IV addresses tests performed with different types of electric fans producing electrical interference near the transmitting devices. Section V covers a test performed with a microwave oven in operation near the transmitting devices and Section VI provides concluding remarks on the tests. The paper closes with references cited.

II. TEST SETUP

A prototype system (Figure 1) was previously developed and is employed here to monitor wireless signals and identify unique edge devices by a naturally occurring fingerprint comprised of polarization characteristics in the wireless analog signals they transmit. This fingerprint is quite unique and very stable for each fixed wireless device. When a set of devices are authenticated based on their fingerprints, a new device entering the area can be identified as an unknown device, even if the perpetrator is using the MAC address and password of an authenticated device to attempt connection with a network. In addition, devices of the same make and model will yield very different fingerprints, making the authentication device specific.

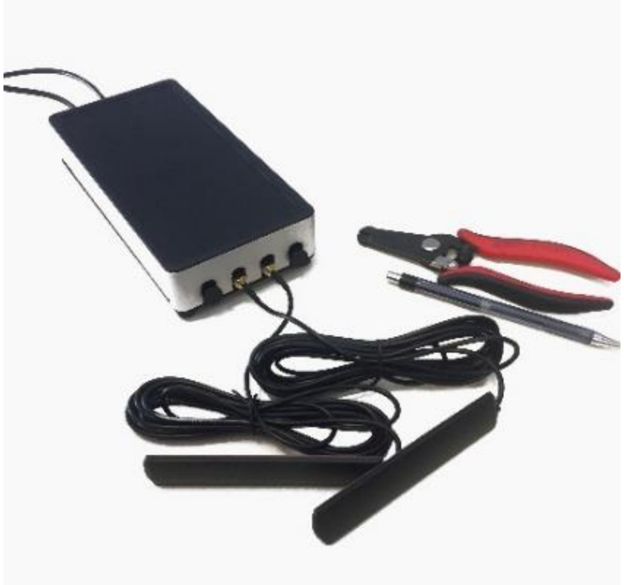

Figure 1. Device under test: a prototype wireless intrusion detection system

Since the fingerprints are naturally occurring, this method of identifying wireless intrusion works with legacy devices, which may have little or no security measures embedded. It also works with any protocol and with any standard for communications, thus making it a potentially desirable mechanism to employ in sensitive industrial environments. However, most industries are electrically noisy and any new security system must be able to operate in an environment with high electrical interference. Thus, a study is needed to ensure its viability for industrial applications.

Electrical interference can raise the noise floor of a received wireless signal. This may result in an otherwise satisfactory signal strength arriving with a low Signal-to-Noise Ratio (SNR), causing problems for some receivers. In addition, electrical interference can result in sporadic increases in energy received, which may appear as new signals, or which may obfuscate desired signals.

To test the prototype under these conditions, a 40-foot by 80-foot metal warehouse was used to simulate a plant environment. It was set up with four sensors each based on the same make and model of microcontroller; in this case, Raspberry Pi 4Bs. The sensors alternately each sent a pair of wireless signals containing data.

Two identical prototypes, each with a unique pair of antennas, were set up to monitor the incoming signals from the sensors for comparison of antenna types. Twelve tests were run using different signal gains and interference sources in varying forms of electrical noise within the field of the transmitting sensors.

The wireless signals were received by each of the prototypes using orthogonally-polarized antennas; one set of RF Elements OARDSBX244 Omni Directional 2.4 GHz, 4dBi antennas and one set of Bestkong Omni WiFi Booster 2.4 GHz 5dBi antennas. The signals were sampled at 20M samples per second and were digitized with a 12-bit Analog-to-Digital Converter. For this test case, they were recorded so they could be analyzed in off-line processing. In actual operating conditions, the inputs would be analyzed as they were received and a decision made as to whether or not they were authenticated, known sources.

Received signals were band-pass filtered and converted to complex baseband in the Universal Software Radio Peripheral [5]. A proprietary pulse detection algorithm was then employed on the baseband signals, and a block of 4096 samples was formed upon detection of a signal. The block was transformed to the frequency domain using a Discrete Fourier Transform (DFT). Further processing, including an algorithm for finding the main spectrum lobe [6], was used to derive a polarization mode dispersion profile across the spectrum of the DFT, eliminating artifacts derived from spectral leakage.

By averaging energy over many symbols within the received packet, one is able to mitigate concerns of incipient deviations, scalloping, unbalanced spectra and other fading phenomena which might influence the calculation of polarization mode dispersion. This computation produces a frequency-dependent fingerprint based on polarization mode dispersion across the signal bandwidth that is quite unique for each sensor. The fingerprint is compared to a bank of collected fingerprints to determine if the received input has been identified previously. This is done through a correlation process, concluding with a number between 0 and 1 that indicates the degree of confidence for each case where the fingerprint of the incoming signal is compared to each known source.

III. Low Relative Signal Strength

A. *Setup*

In this series of experiments, three of the signals received from sensors each had an SNR of 20 dB and a fourth received signal had a 14 dB SNR. These signal strength levels are undesirable in communications and often reflect conditions where the bit-error rates increase to the point where packets fail and must be re-transmitted. Figure 2 depicts a dial which reflects, on average, when wireless communications are good and when they begin to fail. It should be noted that the range from 15 dB to 25 dB is referenced as "poor," indicating that re-transmissions are frequent. Three of the sensors are transmitting signals that are received in this range. The fourth sensor resides in the range below that, denoted as minimum SNR, 10-15 dB. In this range, data may get through only periodically. The degradation is clearly evident in the actual recorded signal spectrum, shown in Figure 3, where the intensity across the spectrum is indicated with a relative intensity, in Volts, on the vertical axis for each frequency bin number on the horizontal axis.

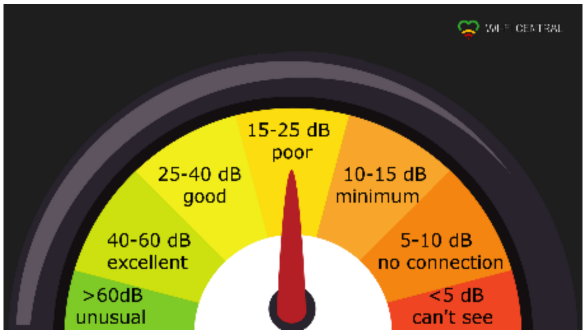

Figure 2. A dial indicating how SNR impacts quality of signal

Typically, one would expect the average spectrum to be approximately symmetric around the center of the chart. Frequency-selectivity due to multipath in the propagation channel can result in signal levels that depend on the frequency, leading to an unbalanced spectrum.

B. *Results*

While this phenomenon and other frequency-selective fading can reduce signal levels and even cause bit-error rates to increase, the effect is not significant enough to negatively impact the ability to correctly match the fingerprint of a received signal with one of a set of known sources. Because the polarization dispersion measurement is averaged over multiple symbols within the packet, transient effects are minimized and an integration gain is achieved.

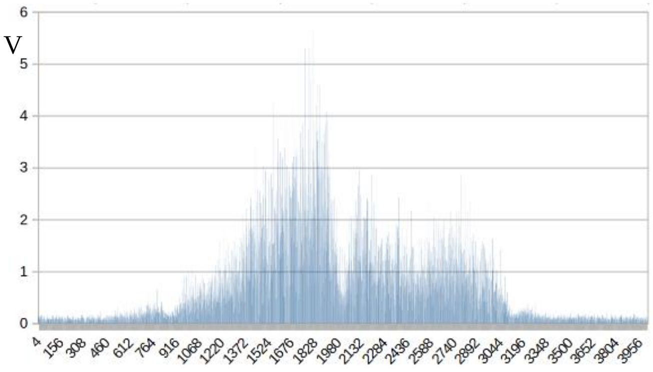

Figure 3. DFT of a signal with low SNR showing non-uniformity in symmetry (test 20210327103909-2G0202e) using relative intensity in Volts on the vertical scale compared at each frequency bin on the horizontal scale

The following chart (Figure 4) is a 'confidence matrix', which is similar to a well-known confusion matrix, except, where a confusion matrix would have known sources along one axis and the same number of unknown sources along a second, the confidence matrix is designed with known sources in columns and each row containing a new, incoming source. In other words, the confidence matrix grows in length with the number of signals received in the test case.

Column A contains the block number of the rising edge of each signal found. The first time a signal is seen, it is not successfully correlated with a known source, since there are no known sources, as yet. Thus, for block 1413 the correlation is 0.38, where 1.00 is a perfect match and 0.00 indicates no correlation.

The confidence matrix shows no sources present in the test other than the four known sensors, shown in columns B, C, D and E. It also shows strong correlation with alternating pairs of signals arriving from each of four sensors. There are no false positive correlations, nor false negative correlations. The number in each cell represents the confidence in making the matching decision. Thus, all positive correlations are above 0.95 indicating that the decisions are made with greater than 95% confidence. In fact, this confidence matrix is from a test with received signals of low SNR and the average positive correlation for this test is 0.97 with a standard deviation of 0.01.

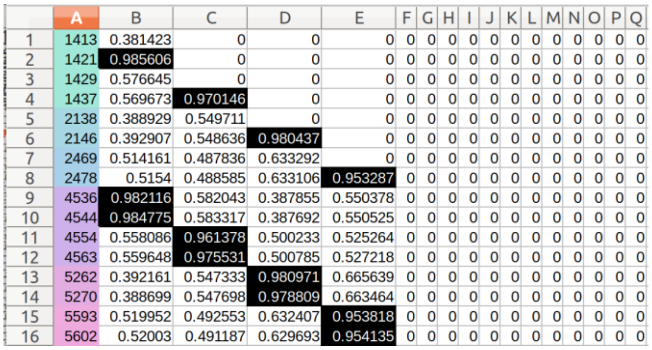

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P	Q
1	1413	0.381423	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2	1421	0.985606	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3	1429	0.576645	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
4	1437	0.569673	0.970146	0	0	0	0	0	0	0	0	0	0	0	0	0	0
5	2138	0.388929	0.549711	0	0	0	0	0	0	0	0	0	0	0	0	0	0
6	2146	0.392907	0.548636	0.980437	0	0	0	0	0	0	0	0	0	0	0	0	0
7	2469	0.514161	0.487836	0.633292	0	0	0	0	0	0	0	0	0	0	0	0	0
8	2478	0.5154	0.488585	0.633106	0.953287	0	0	0	0	0	0	0	0	0	0	0	0
9	4536	0.982116	0.582043	0.387855	0.550378	0	0	0	0	0	0	0	0	0	0	0	0
10	4544	0.984775	0.583317	0.387692	0.550525	0	0	0	0	0	0	0	0	0	0	0	0
11	4554	0.558086	0.961378	0.500233	0.525264	0	0	0	0	0	0	0	0	0	0	0	0
12	4563	0.559648	0.975531	0.500785	0.527218	0	0	0	0	0	0	0	0	0	0	0	0
13	5262	0.392161	0.547333	0.980971	0.665639	0	0	0	0	0	0	0	0	0	0	0	0
14	5270	0.388699	0.547698	0.978809	0.663464	0	0	0	0	0	0	0	0	0	0	0	0
15	5593	0.519952	0.492553	0.632407	0.953818	0	0	0	0	0	0	0	0	0	0	0	0
16	5602	0.52003	0.491187	0.629693	0.954135	0	0	0	0	0	0	0	0	0	0	0	0

Figure 4. Confidence Matrix highlighting positive correlations made each pass between new and known sources, where column A contains the number of the block received; that is, the pass

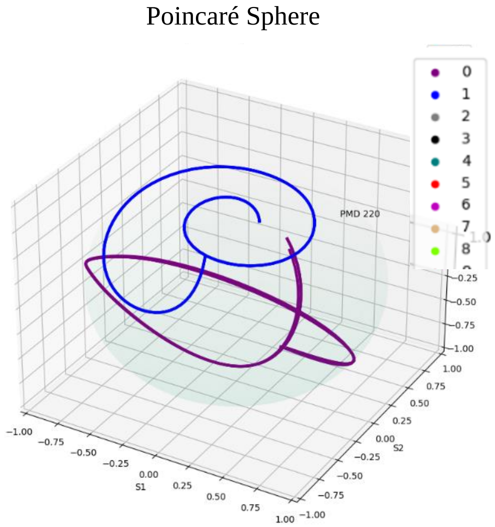

Figure 5. Fingerprints for two devices with low SNR are shown in red and blue curves on the surface of a sphere

This statistic is important when illustrating prototype performance for signals with low SNR. The ability to maintain an average correlation with a confidence factor of 0.97 under low SNR conditions indicates that the fingerprints of wireless edge devices will remain strong even when communications begin to fail.

Polarization measurements from the prototype may be plotted on a spherical coordinate system, called a Poincaré Sphere. Each frequency bin of the DFT contributes a unit vector ending with a single point location on the sphere's surface. In research conducted at the University of Notre Dame, the polarization of a signal has been found to be frequency dependent, leading to a curve, like those shown in Figure 5, as each frequency bin of the DFT is traversed [7].

The fingerprint for a fixed, wireless device may be plotted on a sphere's surface for purposes of visualization, although not necessary for the purpose of correlation. Fingerprints for two separate sensors in the aforementioned chart are color-coded to indicate each device; one, red, and the other, blue. Each point of a fingerprint represents a frequency bin of the DFT. Thus, over the bandwidth of the received signal, a curve meanders around the spherical globe.

Eight tests were run with signals of low SNR. All tests yielded results similar to the test shown above-- there were no false negatives and no false positives.

IV. Electrical Interference

A. Setup

Another set of tests was conducted for conditions involving electrically noisy environments. In these tests, noise-producing equipment was placed near the sensors, one at a time. A table-top rotating fan and a box fan were each used to introduce noise into the environment, one at a time.

The tests involving fans present interesting cases for this technology, since they introduce both electrical interference from the fan motor and motion interference from the rotation of the fan blades within the multi-path.

Both fans were placed, one at a time, on the same counter top as the sensors for this test. This places a fan in close enough proximity to couple electrically with the wireless signals and introduce rotating reflectors in the multi-path environment. This test involved both Bluetooth and Wi-Fi signals, but only the Wi-Fi signals are discussed in this document for simplicity. It suffices to say that no difference was found in the two cases.

B. Results

The DFT of a Wi-Fi signal in this test case appears just as one would expect, with a single main lobe containing a center null. This is the same shape in the frequency response as one would find in a sufficiently strong Wi-Fi signal for good data demodulation. The main lobe is surrounded by two small side lobes, introduced by the finite nature of the DFT. The resulting signal spectrum is shown in Figure 6.

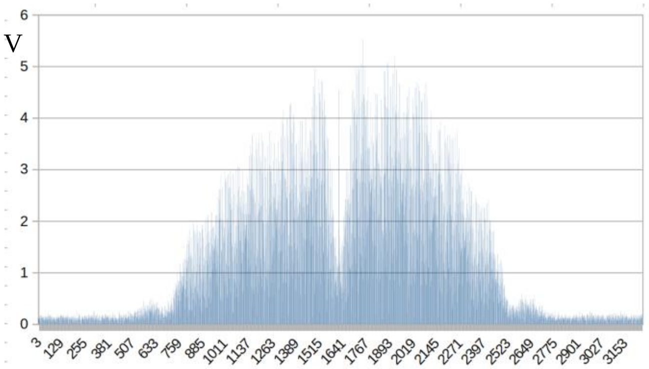

Figure 6. DFT showing relative strength in Volts on the vertical axis and frequency bins on the horizontal axis for a Wi-Fi signal in a test case with a small fan nearby (test I201-F2462-R20-G20-SC16-BS2048-1b)

There were no false positives and no false negatives found in this test case. The average confidence in positive correlations is 98.6% with a standard deviation of 0.01.

As may be seen in Figure 7, which compares a case with no interfering devices, in part A, to the introduced fan, in part B, the noise floor is considerably higher with the fan. Even so, the electrical noise is largely non-polarized and, thus, is for the most part invisible to the algorithms used for fingerprinting. Part C of the figure will be discussed below.

V. MICROWAVE OVEN INTERFERENCE

A. Setup

A test involving interference from a microwave oven is often considered one of the hardest tests to pass in wireless studies. The microwave produces high power signals exactly in the range of frequencies often encountered in the upper ranges of Wi-Fi 2.4 GHz channels. In this test case, we employ channel 11, which is often in the center of such interference.

A microwave oven located near the sensors was turned on for the duration of the test. The result is a series of closely timed pulses which vacillated in amplitude overlaying the sensor signals.

B. Results

Figure 7 shows normal signal amplitude in the time domain in part A, a snapshot of the elevated noise floor from a fan motor in part B, and also a snapshot of microwave background radiation as seen by the receiving antenna on one of the prototypes in part C.

The unusual pulses from the microwave have an effect similar to lowering the received signal SNR by introducing a floor resulting from the presence of the interference. As in the previous tests, the interfering microwave pulses do not seem to significantly influence either the frequency content of individual signals, nor the polarization. Instead, they result in the appearance of a raised noise-plus-interference floor that is not as stable as an environment with no interference.

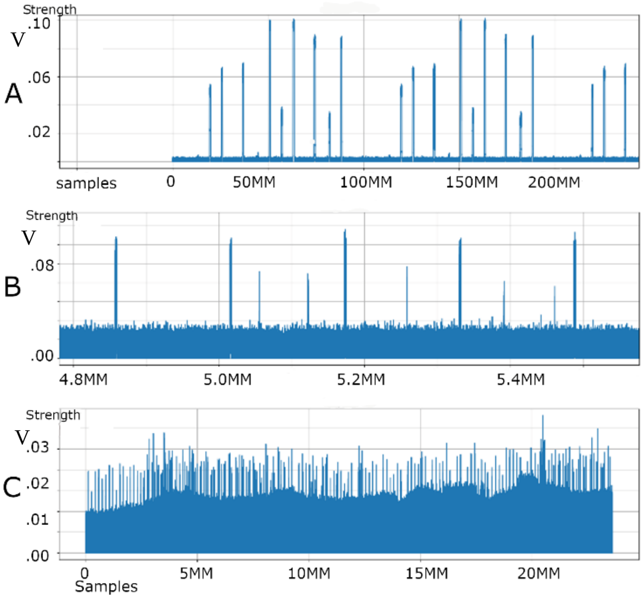

Figure 7. Time domain views of (A) normal sensor activity, (B) fan motor raising the noise floor, and (C) pulses from a microwave raising the noise-interference floor

Taking a closer look at the phenomena, one can see in Figure 8 that the signal spectrum appears as a well-balanced, fairly clean communications signal. Here, we see a main lobe with a null at the center, framed with small side lobes and only a very small amount of deterioration in the right half of the main lobe beginning to form.

Certainly, however, the microwave signal is a major interfering signal and it is clearly seen the in Figure 7 (C) and clearly it lowers the ratio of signal-to-interference-plus-noise of the incoming signals. An examination of the relative intensity (the vertical axis) shows the signal at only half the level of the average intensity of the spectrum in Figure 7.

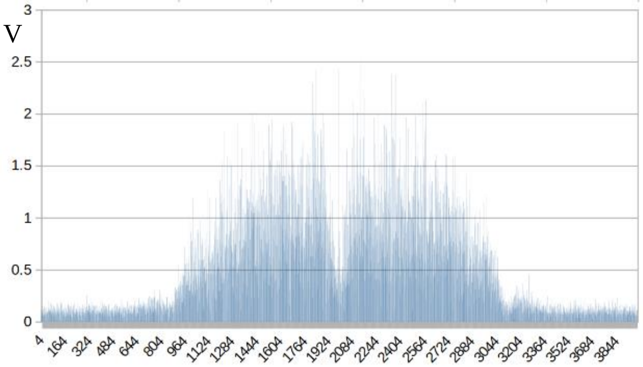

Figure 8. A Discrete Fourier Transform of a signal subject to background microwave interference

If one analyzes a very short period of time, it is possible to see that the background radiation is actually a series of pulses. This may be seen in Figure 9, where a handful of signals fall amidst a continuing series of low amplitude pulses. These do not contain data, of course, but rather are pulses of interfering energy.

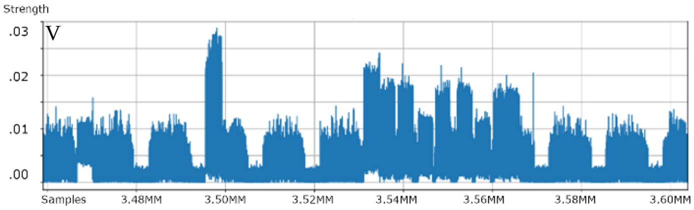

Figure 9. Microwave interference expanded in time shows a series of energy pulses

One might expect that microwave interference would also interfere with the polarization characteristics of the received signals from sensors. However, it was found that the pulses had little effect on the polarization calculation. The average confidence in positive correlation across the test file is 98.0% with a standard deviation of 0.03. Thus, the fingerprints for the identified signals appear to be stable. Fingerprints of the first two devices may be seen in Figure 10. A close examination reveals that the red curve moves very slightly, captured in this image showing the current fingerprint and the previous fingerprint it replaces. Overall, however, the fingerprints for both devices remained fairly stable throughout the test. As in the previous tests, there were no false positives and no false negatives found in this test.

VI. Conclusion

In conclusion, the polarization methodology employed for fingerprinting RF signals from wireless edge devices revealed no false positives and no false negatives in 12 tests directed toward studying low SNR and electrical interference from fans and a microwave oven. The confidence of the positive correlations ranged from 97% to 99%, indicating the methodology is resilient to both conditions of low signal strength and electrical interference. Thus, it may be concluded that the fingerprinting of wireless signals using polarization characteristics is quite robust under conditions of low SNRs and electrical interference. In future work, it is recommended that similar tests be performed to study the effects of motion in the multipath on the methodology.

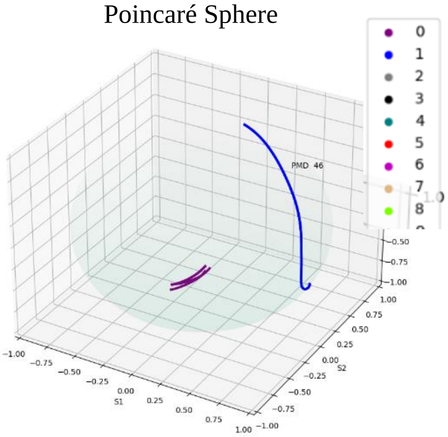

Figure 10. Polarization-based RF fingerprints for two wireless edge devices subject to microwave interference

References

[1] 2019 Global ICS & IIoT Risk Report, published by CyberX, a MicroSoft Azure company, 2019, [retrieved: October, 2021]. Available from https://bit.ly/37wsnix

[2] B. Alotaibi and K. Elleithy, “Rogue access point detection: taxonomy, challenges, and future directions,” Wireless Personal Communications, June 11, 2016, pp. 1261-1290, DOI: 10.1007/s11277-016-3390-x.

[3] R. P. Heller, T. G. Pratt, J. Loof and E. Jesse, “RF biometric for wireless devices,” Proceedings of the Future Technologies Conference (FTC) 2018. FTC 2018. Advances in Intelligent Systems and Computing, vol 881. Springer Nature Switzerland AG, Cham., October 2018. https://doi.org/10.1007/978-3-030-02683-7_65

[4] P. Heller, “Wireless frequency data manipulation for embedded databases use in cybersecurity applications,” International Conference on Digital Communications, ICDT, April 18, 2021, pp. 36-42, ISBN: 978-1-61208-835-8.

[5] J. Petrich, VHF-UHF-microwave SDR transceiver on the air, American Radio Relay League Northwestern Division Convention, SEA-PAC 2017, June 3, 2017. [retrieved: October, 2021] Available from: https://seapac.org/seminars/2017/SEA-PAC2017-uhf-sdr-transceiver.pdf

[6] P. Heller, “Radio frequency fingerprinting with polarizaton mode dispersion,” a tutorial, Nexcom, Porto, Portugal, April 18, 2021. Available from: https://youtu.be/zE2whBIoaAI.

[7] T. G. Pratt and R. D. Kossler, “Input-to-output instantaneous polarizaton characterization,” IEEE Transactions on Antennas and Propogation, Vol. 67, No. 3, pp. 1804-1818, March 2019.