AFRL-RY-WP-TP-2019-0074

A REVIEW OF ANOMALY DETECTION TECHNIQUES LEVERAGING SIDE-CHANNEL EMISSIONS (Preprint)

Igor Ternovskiy

Antennas and Electromagnetics Technology Branch
Multispectral Sensing & Detection Division

Devin Spatz and Devin Smarra

Defense Engineering Corporation

APRIL 2019
Final Report

Approved for public release; distribution is unlimited.

See additional restrictions described on inside pages

STINFO COPY

AIR FORCE RESEARCH LABORATORY
SENSORS DIRECTORATE
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7320
AIR FORCE MATERIEL COMMAND
UNITED STATES AIR FORCE

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. **PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS**.

1. REPORT DATE *(DD-MM-YY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
April 2019	Conference Paper Preprint	15 April 2019 – 15 April 2019

4. TITLE AND SUBTITLE
A REVIEW OF ANOMALY DETECTION TECHNIQUES LEVERAGING SIDE-CHANNEL EMISSIONS (Preprint)

5a. CONTRACT NUMBER
N/A

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
N/A

6. AUTHOR(S)
Igor Ternovskiy (AFRL/RYMH)
Devin Spatz and Devin Smarra (Defense Engineering Corporation)

5d. PROJECT NUMBER
N/A

5e. TASK NUMBER
N/A

5f. WORK UNIT NUMBER
N/A

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Sensing Management Branch, Multi-Domain Sensing
Autonomy Division
Air Force Research Laboratory, Sensors Directorate\
Wright-Patterson Air Force Base, OH 45433-7320
Air Force Materiel Command, United States Air Force

Defense Engineering Corporation

8. PERFORMING ORGANIZATION REPORT NUMBER
AFRL-RY-WP-TP-2019-0074

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Air Force Research Laboratory
Sensors Directorate
Wright-Patterson Air Force Base, OH 45433-7320
Air Force Materiel Command
United States Air Force

10. SPONSORING/MONITORING AGENCY ACRONYM(S)
AFRL/RYMH

11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S)
AFRL-RY-WP-TP-2019-0074

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
PAO case number 88ABW-2019-1670, Clearance Date 11 April 2019. To be presented at SPIE DCS Defense + Commercial Sensing: Cyber Sensing 2019, Baltimore, MD. This work was funded in whole or in part by DARPA LADS. The U.S. Government is joint author of this work and has the right to use, modify, reproduce, release, perform, display, or disclose the work. Report contains color.

14. ABSTRACT
The low-cost, flexible nature of Internet-of-Things (IoT) hardware has resulted in widespread usage in a variety of applications from smart-home systems to industrial process-regulation controllers. As the number of network-connected IoT devices has proliferated, they have become increasingly likely to be the target of widespread cyber-attacks. Since these devices are often low-resource, embedded or bare metal systems, conventional profiling techniques used by Personal Computers (PCs) and workstations have become highly impractical means for security. As a result, an IoT device could provide intruders with an unprotected backdoor into a network. Effectively protecting IoT hardware requires that alternative security protocols be developed and utilized to protect the IoT and the networks they are integrated with. One potential way of improving the security of IoT devices is by monitoring their side-channel emissions to observe device behavior. As these devices operate, they will produce multi-spectral phenomenon, or side-channel emissions, that correlate with program execution. By combining spectral analysis techniques with powerful machine learning algorithms, side-channel emissions can be utilized to bolster IoT device security and deny an intruder access to the network. This paper will review current state-of-the-art techniques used to monitor and classify the behavior of IoT devices. The paper will conclude by discussing several real-world applications presented in literature that have been shown to benefit from these techniques.

15. SUBJECT TERMS
anomaly detection, side-channel analysis, feature extraction, machine learning, deep learning, classification, graph theory

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT:	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON (Monitor)
a. REPORT	b. ABSTRACT	c. THIS PAGE			Igor Ternovskiy
Unclassified	Unclassified	Unclassified	SAR	10	19b. TELEPHONE NUMBER *(Include Area Code)* N/A

A review of anomaly detection techniques leveraging side-channel emissions

Devin Spatz[a,b], Devin Smarra[a,b], and Igor Ternovskiy[b]

[a]Defense Engineering Corp. (United States)
[b]Air Force Research Laboratory (United States)

ABSTRACT

The low-cost, flexible nature of Internet-of-Things (IoT) hardware has resulted in widespread usage in a variety of applications from smart-home systems to industrial process-regulation controllers. As the number of network-connected IoT devices has proliferated, they have become increasingly likely to be the target of widespread cyber-attacks. Since these devices are often low-resource, embedded or bare metal systems, conventional profiling techniques used by Personal Computers (PCs) and workstations have become highly impractical means for security. As a result, an IoT device could provide intruders with an unprotected backdoor into a network. Effectively protecting IoT hardware requires that alternative security protocols be developed and utilized to protect the IoT and the networks they are integrated with. One potential way of improving the security of IoT devices is by monitoring their side-channel emissions to observe device behavior. As these devices operate, they will produce multi-spectral phenomenon, or side-channel emissions, that correlate with program execution. By combining spectral analysis techniques with powerful machine learning algorithms, side-channel emissions can be utilized to bolster IoT device security and deny an intruder access to the network. This paper will review current state-of-the-art techniques used to monitor and classify the behavior of IoT devices. The paper will conclude by discussing several real-world applications presented in literature that have been shown to benefit from these techniques.

Keywords: Anomaly Detection, Side-Channel Analysis, Feature Extraction, Machine Learning, Deep Learning, Classification, Graph Theory

1. INTRODUCTION

The Internet of Things (IoT) consists of a rapidly increasing number of low-cost, low-resource computing devices that leverage widespread connectivity for automation across consumer, industrial, and government sectors. Aspects such as cost and ease of use have made IoT devices popular among amateurs and professionals alike. While IoT devices enable connectivity and automation in many systems, they also present major security concerns due to their low-resource nature. Conventional security techniques, such as profiling, are generally executed directly on system hardware in order to provide a computer with crucial information about its own performance. However, the characteristically low-resource nature of IoT devices has made it so conventional security techniques are incompatible with their architecture.[1] Being unable to monitor device performance internally leaves IoT devices and their networks vulnerable to attacks. Therefore, as the number of cyber-attacks on industry and government resources has increased in recent years, it should come as no surprise that IoT devices are targets. The inability of IoT devices to integrate cyber-defense into their firmware leaves them relatively unprotected backdoors to their respective networks.[1]

An alternative approach to conventional security techniques is cyber-physical sensing. Cyber-physical sensing leverages unintended device emissions to determine which operations are executing on a device. Often, this approach requires a suite of sensors to monitor device emissions. These emissions, which correlate with the behavior of the program being executed, can provide valuable markers to classifying a device's behavior.[2] For example, every bit-flip, switch of a transistor, or change in electrical current path can produce unique acoustic,

Further author information: (Send correspondence to Igor Ternovskiy)
Igor Ternovskiy: E-mail: igor.ternovskiy@us.af.mil

Approved for public release; distribution is unlimited

electrical, and thermal signatures.[3] The use of sensor systems with fine-tuned sensing parameters, such as bandwidth and imaging resolution, enables the detection of even the most minute change in these emissions. Strong signal processing and machine/deep-learning algorithms are generally applied in order to classify the data where the resulting classification can indicate whether or not a program has deviated from its intended mode of operation. As such, cyber-physical sensing enables remote monitoring and detection of anomalous behavior, lending themselves to be potentially viable solutions to the security challenges of IoT devices.

While many of the approaches to cyber-physical sensing utilize the same basic philosophy; measure, collect, process, and classify, they differ greatly in their execution and architecture. The remainder of this paper will be spent investigating the various ways that researchers are using cyber-physical sensing to enhance the security of IoT devices utilizing the electromagnetic (EM) sensing modality.

2. SIDE-CHANNEL EMISSION LOCALIZATION & MEASUREMENT

Side-channel emissions from IoT devices serve as a source of information about device operation when signals can be effectively identified and measured. A primary challenge involves determining a measurement probe position which provides the strongest RF side-channel signals while mitigating interference from external noise and interference.[4] Effective techniques in localization and measurement together can be utilized to produce the highest quality EM signals for IoT behavior classification.

2.1 LOCALIZATION

Due to the unintended nature of EM side-channel emissions, there are no predictable locations for ideal signal measurement either on-board or on-chip in IoT devices. Traditional EM side-channel localization techniques are time-intensive, requiring a full grid search over the surface of an IoT device to localize the strongest EM signal sources.[4] In order to reduce localization times, Werner et al. developed an efficient localization algorithm to reduce the overall number of measurements required for magnetic source localization.[4,5] The proposed algorithm solves a forward-backward optimization problem using a magnetic field model and measurements taken around the edges of the device to identify one or more sources.[4] Experimental validation of the algorithm demonstrates the ability to successfully localize on-board and on-chip sources for FPGA and IoT devices when one or more sources are present.[5]

Although raw EM signal strength is an effective metric for localizing signal sources, simple power metrics such as the root-mean-squared (RMS) amplitude of signals can be biased by background signals and noise.[6] An alternative approach to determining ideal measurement locations on-board IoT devices is classifier performance evaluation. Riley et al. approaches source localization by performing a grid search above the processor and training a Support-Vector Machine (SVM) on the collected traces from each grid location to find the location with the lowest probability of misclassification. This method is effective in localizing sources with minimal interference as the SVM error metric will provide degraded results when high levels of noise and interference are present.[6–8]

2.2 SIDE-CHANNEL MEASUREMENT

A primary challenge in EM side-channel measurement is the inherent lack of signal strength. Typically, the Signal-to-Noise Ratio (SNR) is used to quantify signal strength, but the traditional definition does not account for sources of noise other than thermal noise. Juyal et al. establishes a definition of SNR for EM side-channels by considering the noise to be a combination of constant thermal noise and noise generated by the processor which reduces in strength following the propagation equation. Experimental measurements of EM side-channel emissions are shown to validate this SNR model providing insight on how distance effects side-channel measurement quality.[9] Furthermore, Zajic et al. investigates propagation mechanisms of EM side-channels to model path loss behavior in the near-field and far-field. Experimental measurements of Amplitude Modulated (AM) carriers along with upper and lower side-bands are found to be in agreement with the proposed propagation loss model.[10]

Approaching side-channel measurement with a better understanding of specific SNR characteristics, antenna design and characterization are considered to increase measurement SNR.[9,11–13] Graham et al. establishes

Approved for public release; distribution is unlimited

a characterization method for near-field probes as antenna characteristics impose the greatest limitations on signal sensitivity and frequency response for EM side-channel measurement. Using the proposed measurement method, multiple near-field probes are characterized to determine peak operating frequencies.[13] Juyal et al. approaches the SNR problem through directive antenna design as a means to increase signal strength and mitigate interference. A circular disc antenna design is presented which achieves a peak gain of 19 dBi while operating in the TM12 mode.[9] The antenna design problem for side-channel measurement is further considered by Adibelli et al. in which a non-uniform helical antenna design is realized. This design is shown to provide 2-3 dB more directivity than optimal uniform helix designs of the same length.[11,12] These improvements in antenna gain for traditional antenna designs increase measurement capability while maintaining antenna size. Together, these advancements in antenna characterization and design permit more-robust measurement of EM side-channels which, in turn, aid in the feature identification and extraction process.

3. FEATURE IDENTIFICATION & EXTRACTION

The use of machine learning algorithms is essential for processing the large volumes of side-channel emissions data produced by IoT devices. In order to utilize these algorithms effectively, refined feature identification and extraction procedures are crucial. Furthermore, in instances where program binaries are available, it is possible to gain additional insight about features and, potentially, modify programs in a manner conducive to feature extraction.

3.1 Feature Identification

A key component of side-channel analysis is to identify features which carry information about the underlying source phenomena. From an information theory standpoint, not all side-channel emissions carry information about the internal processor states. Yilmaz et al. examines processor instruction execution to quantify information and channel capacity.[14–16] Through analysis of a worst case scenario where the program is specifically designed to leak information about the device behavior, Yilmaz et al. determined the information capacity to be 1.405 bits/symbol.[16] Yilmaz et al. also approached side-channel analysis through a communications perspective by modelling the side-channel emissions as a Pulse Amplitude Modulation (PAM) signal. From this model, the channel Power Spectral Density (PSD) and Bit Error Rate (BER) bounds are established for a covert side-channel with jitter.[17]

Examination of the structure of side-channel signals in the frequency domain provides insight into clock frequencies/harmonics as well as the modulation of events onto such carriers. Sehatbakhsh et al. identifies frequencies relative to the device clock frequency and harmonics which correspond to periodic activity in programs (i.e. loops). By determining the per-iteration execution time for each loop for a given processor architecture and instruction set, it is possible to estimate the offset frequency to be the frequency corresponding to the execution period.[18] Prvulovic et al. proposes a method for finding Amplitude-Modulated (AM) and Frequency-Modulated (FM) EM side-channel emissions. The assignment of a confidence level and estimated Signal-to-Noise (SNR) ratio to each identified AM or FM feature serves as a metric for choosing potential features and disregarding others.[19] Wang et al. builds upon this work and presents an algorithm for automatically finding carriers of AM emanations in the EM spectrum.[20]

3.2 Feature Extraction

Throughout literature, selection of signals in the frequency domain is commonly used as in input to machine learning classifiers. This is often due to that fact that feature extraction in the time domain requires precise alignment of traces such that features remain consistent across all measurements.[6] Agrawal et al. searches the frequency spectrum to determine bandwidths over which device emissions are unique from those of the environment. Using the optimum frequency band, Discrete Fourier Transform (DFT) based features are extracted via the Short Time Fourier Transform (STFT) with a sliding window.[21] Nazari et al. approaches extraction in a similar manner by creating Short-Time Spectrum (STS) signatures of the EM emissions spectrum by applying the STFT to windowed sequences of the continuous time signal. This method allows for the capture of loop executions and transitions between loops.[22] Boggs et al. utilizes all frequency bins over a measured bandwidth

Approved for public release; distribution is unlimited

to extract signatures representative of different processor states. Signatures are determined to be unique through a similarity metric and signatures with high similarity are merged to produce robust signatures.[23]

A common challenge in feature extraction is reducing a feature vector to a manageable number of features for effective application of machine learning. Usage of all bins in the FFT spectrum as a source of features has several problems including increased computational complexity and the risk of overtraining classifiers.[21] Spatz et al. examines how a feature vector of FFT based features can be reduced using various measurement metrics. Selection of FFT based features using feature mean and variance is found to be effective in reducing a feature vector while still maintaining good classification training and testing performance.[24] Agrawal et al. utilizes mutual information as a metric for each of the FFT based features to select a reduced feature set that contains nearly the same information as the underlying feature vector.[21]

3.3 Firmware Analysis & Boosting

In instances where program binaries and device architectures are known, it is possible to establish ground truth data for expected program execution.[3,21] Graham et al. emulates the Arduino Uno device architecture to produce ground truth data for when specified program "blocks" will execute and the expected duration of each block. The time information established for each block enables ground truth states to be mapped to emanations occurring within the specified time window.[3] Agrawal et a. takes a similar approach to device characterization through which micro-benchmarks are derived from device firmware. Micro-benchmarks, which consist of small programs containing instructions from segments of known code, are selected for instruction segments expected to generate "loud" analog emissions. These micro-benchmarks contain labels which together can be used for differentiating between program states.[21] Dey et al. uses a cycle-accurate architectural simulator to create ground-truth data for processor stalls and memory- access events. Collected near-field EM emanations for these events were found to correspond to the simulated power consumption.[25] In each of these instances, access to program binaries along with device architecture emulation capabilities enables the creation of ground truth data. Unfortunately, this capability is limited by the speed and complexity of device architectures which can be emulated.[3]

An alternate approach is taken by Boggs et al. in which start and end times of events are recorded by the remote host interacting with a networked device. The log of these events provides ground truth data for when certain events occur and the expected durations of these events.[23] Finally, in instances where program binaries are available for modification, processor states can be made more observable through a technique called "boosting". Agrawal et al. discusses how additional code can be inserted into program binaries such that stronger known emissions will appear at specific points during program execution.[21] This type of technique enables easier state identification and tracking. All of these techniques leverage available resources to make program states more observable for extraction and classification.

4. MACHINE & DEEP LEARNING APPROACHES

Machine & deep learning provide much of the intelligence behind anomaly detection. Utilizing the features discussed in the previous section, a number of techniques are used across supervised-, unsupervised-, and deep-learning to classify IoT device states from EM side-channel emissions.

4.1 Supervised Machine Learning

Techniques in supervised Machine Learning (ML) commonly serve as the basis for classification of side-channel emissions. Depending on measurement fidelity, supervised ML is used to classify behavior at the instruction level,[26] block level,[3] or program/firmware level.[27] At the instruction level, Riley et al. extracts the signatures of each processor cycle and performs background carrier subtraction to obtain Program Counter (PC) and Instruction Register (IR) waveforms that occur for bit-flips. The resulting waveforms are used to train an SVM classifier on the hamming distance of the Stack Pointer (SP) and IR.[26] Graham et al. utilizes processor emulation to create blocks of varying length from program code from which corresponding side-channel emissions are used to train a block classifier. The classifier is built up of multiple binary classifiers which are able to distinguish between device states at the block level.[3] Finally, Fuller et al. trains various supervised ML algorithms on EM side-channel emissions from an FPGA program to distinguish between different bitstreams which contain differences in logic and placement of gates within the FPGA.[27]

Approved for public release; distribution is unlimited

Across the literature, SVM classifiers are found to be an effective supervised ML algorithm for classification of EM side-channel emanations. In particular, the linear kernel function for the SVM classifier was found to deliver good performance across datasets.[8,21] Another important aspect of data collection involves preventing spurious noise sources from becoming trained features in a classifier. Riley et al. approaches this aspect of training by interleaving the collection of training data to prevent machine learning algorithms from over-training on background signals.[7]

4.2 Unsupervised Machine Learning

A major challenge in applying supervised machine learning algorithms is that the input dataset must be labeled. This creates a problem when working with side-channel emissions data as training data with ground truth labels is not easily obtainable. Unsupervised machine learning techniques such as clustering algorithms prove useful in the task of automatically labeling training data. Agrawal et al. utilizes unsupervised machine learning through a combination of Principle Component Analysis (PCA) and the Density-based Spatial Clustering of Applications with Noise (DBSCAN) algorithms to cluster emissions together into pseudo-states of device operation. The resulting clusters are used as the labels for training supervised machine learning algorithms on the underlying data.[21]

4.3 Deep Learning

Traditional machine learning techniques can provide high classification accuracy for side-channel signals when feature vectors are well engineered. In practice, these techniques fall short as environmental changes (noise) and device-to-device variation challenge even the strongest features. An alternate approach is to utilize deep learning to learn features directly from the data. Zhou et al. trained Multi-Layer Perceptron (MLP) and Long Short-Term Memory (LTSM) networks using raw data represented in the time domain or frequency domain, respectively. The classifiers perform effectively on raw data achieving near 100% accuracy on binary classification problems and 90% accuracy on multi-class classification problems.[2] The effectiveness of deep learning extends beyond classification as it can be used as a powerful tool for detection of features. Wang et al. proposes the utilization of deep learning to remove false positive detections of AM modulated spectra. In this instance, deep learning algorithms can be trained to identify which frequencies are AM spectra and which are not.[20]

5. ANOMALY DETECTION & CONTROL FLOW TRACKING

5.1 Detection

Following the application of machine learning techniques, it is necessary to detect deviations from normal, trained program behavior. Statistical analysis is a commonly used technique for detecting anomalies in RF-based emanations.[21,22] Agrawal et al. approaches the detection problem by computing the Mahalanobis distance between training and testing statistical distributions to declare known or unknown code.[21] Nazari et al. also approaches the detection problem from a statistical standpoint by applying the Kolmogorov-Smirnov test (K-S test) to detect deviations in the testing behavior from the training behavior.[22] These statistical detection methods work well without known information about devices, but lack the ability to detect out-of-order state execution.[21]

A complimentary approach to anomaly detection is to monitor the order of state occurrences to determine if programs are executing as expected. Agrawal et al. proposes an n-gram detector to monitor label execution order. N-grams, or contiguous sequences of n states, can be used to represent short sequences of state order execution. The n-gram detector will indicate when sequences of testing states occur in an order that is not present in the training data.[21] Furthermore, Agrawal et al. discusses a Nondeterministic Finite Automaton (NFA) which builds a model of all the valid state execution paths. As state labels are predicated, validation can be performed against the trained NFA model to determine if execution is occurring outside of the trained valid paths.[21]

While most detection methods focus on order and distributions of device state outputs, Zhou et al. utilizes an autoencoder to perform anomaly detection directly on testing data sequences. During training of the autoencoder network, a sequence of samples is fed into a 5-layer auto encoder. The input and output layers of the network

Approved for public release; distribution is unlimited

have the same number of nodes while a bottleneck is introduced in the middle of the network so that a low-dimensional representation of the data must be learned. Due to this bottleneck, the network learns how to reconstruct the original signal from the low-dimensionality data, and an error metric is used to measure how well reconstruction is performed. Detection can be performed by sending testing data through the network and flagging data sequences with high reconstruction error as these sequences likely have differing data properties.[2]

5.2 Control Flow Tracking

Effective detection of increasingly sophisticated malware requires both the identification of device states and state execution order.[21] Program analysis at this granularity requires the creation of a Control Flow Graph (CFG) which represents all program states and valid transitions between states that may occur between them during normal program execution. In particular, a colored graph model is considered as measured emissions spectra are mapped to graph nodes or edges as colors which are not necessarily unique.[28] Within CFG theory, there exist multiple observability classes for colored graph models including Hidden Markov Model (HMM), trackable, unifilar, and observable. Effective CFG tracking requires an observable graph in which the next state in the graph can be uniquely determined by a number of previous, consecutive states.[29]

The prospect of modifying CFG structure to improve EM-based emission tracking performance of processor behavior has been considered recently in literature.[28,29] Chilenski et al. proposes the addition of transition nodes to CFGs to improve tracking of RF-based processor behavior by reducing tracking performance errors. By adding transition nodes to the CFG through changing or adding code blocks which are highly observable, multiple similar branches or sections of a CFG can be uniquely distinguished between in order to identify current positioning in a CFG.[29] Furthermore, a new framework is proposed which can be used to reason how to place transition nodes to successfully mitigate CFG structures which prevent observability.[28]

Various forms of CFG tracking are used throughout literature to perform anomaly detection. Graham et al. describes a block classifier which consists of multiple binary classifiers, each trained on the potential blocks that a previous known block can transition to.[3] Khan takes a similar approach in creating a dictionary of EM patterns during training which correspond to all possible combinations of branch outcomes that can occur in a section of a program.[30] Finally, the NFA detector proposed by Agrawal et al. builds a CFG from training data to determine when tested device-states, execute at an invalid time.[21] Each of these applications demonstrate how CFG tracking can be used to aid in the identification of device states and detection of devices operating anomalously.

6. APPLICATIONS

Recent advances in anomaly detection research have many applications in securing the IoT through wireless monitoring. Sehatbakhsh et al. extended anomaly detection research to medical applications for monitoring embedded medical devices. The proposed SYNDROME framework monitors the EM signals from medical devices and alerts users of statistical deviations from a trained, malware- free model.[1] SYNDROME was demonstrated on an automated syringe pump system in which malware could be detected when being loaded onto the syringe pump controller. The low-latency detection of the change in device behavior triggered an external shutoff to the syringe pump to prevent any operational errors.[31]

Another more-common concern for the IoT is widespread infection of devices with a bot-net or Ransomware. Sehatbakhsh et al. demonstrates how an embedded IoT device running the Debian Linux operating system can be monitored with commercial off-the-shelf (COTS) hardware to detect changes in the EM spectrum when an IoT device becomes infected.[32] This change- detection capability is a flexible approach to device monitoring as the nature of the malicious code does not need to be known for a general change in operation to be identified and reported. One final selected application from literature is the application of these techniques to RSA key extraction. Alam et al. presents an approach for utilizing side-channel EM emissions to recover the RSA exponent bits. A mitigation approach is proposed to fix this vulnerability in which bits are obtained in groups rather than one bit at time.[33]

ACKNOWLEDGMENTS

The authors would like to acknowledge DARPA LADS for funding this continuing research.

Approved for public release; distribution is unlimited

REFERENCES

[1] Sehatbakhsh, N., Alam, M., Nazari, A., Zajic, A., and Prvulovic, M., "Syndrome: Spectral analysis for anomaly detection on medical iot and embedded devices," in [*2018 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*], 1–8 (April 2018).

[2] Wang, X., Zhou, Q., Harer, J., Brown, G., Qiu, S., Dou, Z., Wang, J., Hinton, A., Gonzalez, C. A., and Chin, P., "Deep learning-based classification and anomaly detection of side-channel signals," *Proc.SPIE* **10630**, 10630 – 10630 – 8 (2018).

[3] Graham, J. T., Riley, R., Baldwin, R., and Fisher, A., "Block-level algorithm classification based on rf side-channel," *Proc.SPIE* **10630**, 10630 – 10630 – 7 (2018).

[4] Werner, F., Chu, D. A., Djordjević, A. R., Olćan, D. I., Prvulovic, M., and Zajić, A., "A method for efficient localization of magnetic field sources excited by execution of instructions in a processor," *IEEE Transactions on Electromagnetic Compatibility* **60**, 613–622 (June 2018).

[5] Werner, F. T., Djordjević, A. R., Olćan, D. I., Prvulovic, M., and Zajić, A., "Experimental validation of localization method for finding magnetic sources on iot devices," in [*2018 International Symposium on Electromagnetic Compatibility (EMC EUROPE)*], 413–418 (Aug 2018).

[6] Riley, R. A., Graham, J. T., Fuller, R. M., Baldwin, R. O., and Sampathkumar, A., "Extraction and validation of algorithms based on analog side-channels," *Proc.SPIE* **10185**, 10185 – 10185 – 15 (2017).

[7] Riley, R. A., Graham, J. T., Swartz, P. D., Baldwin, R. O., and Fisher, A., "Classification of low-snr side channels," *Proc.SPIE* **10630**, 10630 – 10630 – 7 (2018).

[8] Riley, R., Graham, J. T., Fuller, R. M., Baldwin, R. O., and Fisher, A., "Generalization of algorithm recognition in rf side channels between devices," *Proc.SPIE* **10630**, 10630 – 10630 – 7 (2018).

[9] Juyal, P., Adibelli, S., Sehatbakhsh, N., and Zajic, A., "A directive antenna based on conducting disks for detecting unintentional em emissions at large distances," *IEEE Transactions on Antennas and Propagation* **66**, 6751–6761 (Dec 2018).

[10] Zajic, A., Prvulovic, M., and Chu, D., "Path loss prediction for electromagnetic side-channel signals," in [*2017 11th European Conference on Antennas and Propagation (EUCAP)*], 3877–3881 (March 2017).

[11] Adibelli, S., Golubović, R., Djordjević, A., Olćan, D., and Zajić, A., "Design and fabrication of nonuniform helical antennas for detection of side-channel attacks in computer systems," in [*12th European Conference on Antennas and Propagation (EuCAP 2018)*], 1–5 (April 2018).

[12] Dinkic, J., Olcan, D., Djordjevic, A., and Zajic, A., "Comparison of optimization approaches for designing nonuniform helical antennas," 1581–1582 (07 2018).

[13] Graham, J. T., Baldwin, R. O., and Sampathkumar, A., "Characterization of riscure 1-ghz low sensitivity probe for side channel analysis (sca)," *Proc.SPIE* **10185**, 10185 – 10185 – 10 (2017).

[14] Yilmaz, B. B., Prvulovic, M., and Zajić, A., "Capacity of deliberate side-channels created by software activities," in [*MILCOM 2018 - 2018 IEEE Military Communications Conference (MILCOM)*], 237–242 (Oct 2018).

[15] Yilmaz, B. B., Callan, R. L., Prvulovic, M., and Zajić, A., "Capacity of the em covert/side-channel created by the execution of instructions in a processor," *IEEE Transactions on Information Forensics and Security* **13**, 605–620 (March 2018).

[16] Yilmaz, B. B., Callan, R., Prvulovic, M., and Zajić, A., "Quantifying information leakage in a processor caused by the execution of instructions," in [*MILCOM 2017 - 2017 IEEE Military Communications Conference (MILCOM)*], 255–260 (Oct 2017).

[17] Yilmaz, B. B., Zajić, A., and Prvulovic, M., "Modelling jitter in wireless channel created by processor-memory activity," in [*2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*], 2037–2041 (April 2018).

[18] Sehatbakhsh, N., Nazari, A., Zajic, A., and Prvulovic, M., "Spectral profiling: Observer-effect-free profiling by monitoring em emanations," in [*2016 49th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*], 1–11 (Oct 2016).

[19] Prvulovic, M., Zajić, A., Callan, R. L., and Wang, C. J., "A method for finding frequency-modulated and amplitude-modulated electromagnetic emanations in computer systems," *IEEE Transactions on Electromagnetic Compatibility* **59**, 34–42 (Feb 2017).

Approved for public release; distribution is unlimited

[20] Wang, C., Callan, R., Zajic, A., and Prvulovic, M., "An algorithm for finding carriers of amplitude-modulated electromagnetic emanations in computer systems," in [*2016 10th European Conference on Antennas and Propagation (EuCAP)*], 1–5 (April 2016).

[21] Agrawal, H., Chen, R., Hollingsworth, J. K., Hung, C., Izmailov, R., Koshy, J., Liberti, J., Mesterharm, C., Morman, J., Panagos, T., Pucci, M., Sebüktekin, I., Alexander, S., and Tsang, S., "Casper: an efficient approach to detect anomalous code execution from unintended electronic device emissions," *Proc.SPIE* **10630**, 10630 – 10630 – 20 (2018).

[22] Nazari, A., Sehatbakhsh, N., Alam, M., Zajic, A., and Prvulovic, M., "Eddie: Em-based detection of deviations in program execution," in [*2017 ACM/IEEE 44th Annual International Symposium on Computer Architecture (ISCA)*], 333–346 (June 2017).

[23] Boggs, N., Chau, J. C., and Cui, A., "Utilizing electromagnetic emanations for out-of-band detection of unknown attack code in a programmable logic controller," *Proc.SPIE* **10630**, 10630 – 10630 – 16 (2018).

[24] Spatz, D. W., Smarra, D. A., and Ternovskiy, I. V., "Preliminary classification results of rf emission based feature extraction in internet of things devices," *Proc.SPIE* **10630**, 10630 – 10630 – 9 (2018).

[25] Dey, M., Nazari, A., Zajic, A., and Prvulovic, M., "Emprof: Memory profiling via em-emanation in iot and hand-held devices," in [*2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*], 881–893 (Oct 2018).

[26] Riley, R. A., Graham, J. T., Baldwin, R. O., and Fisher, A., "Register hamming distance from side channels," *Proc.SPIE* **10630**, 10630 – 10630 – 8 (2018).

[27] Fuller, R. M., Riley, R. A., and Graham, J. T., "Exploiting side-channel emissions to detect changes in fpga firmware," *Proc.SPIE* **10630**, 10630 – 10630 – 11 (2018).

[28] Chilenski, M., Cybenko, G., Dekine, I., Kumar, P., and Raz, G., "Observability Properties of Colored Graphs," *arXiv e-prints* (Nov. 2018).

[29] Chilenski, M., Cybenko, G., Dekine, I., Kumar, P., and Raz, G., "Control flow graph modifications for improved rf-based processor tracking performance," *Proc.SPIE* **10630**, 10630 – 10630 – 13 (2018).

[30] Khan, H. A., Alam, M., Zajic, A., and Prvulovic, M., "Detailed tracking of program control flow using analog side-channel signals: a promise for iot malware detection and a threat for many cryptographic implementations," *Proc.SPIE* **10630**, 10630 – 10630 – 14 (2018).

[31] Sehatbakhsh, N., Hong, H., Lazar, B., Johnson-Smith, B., Yilmaz, O., Alam, M., Nazari, A., Zajic, A., and Prvulovic, M., "Syndrome: Spectral analysis for anomaly detection on medical iot and embedded devices-experimental demonstration," (05 2017). Hardware Demo at IEEE International Symposium on Hardware Oriented Security and Trust (HOST).

[32] Sehatbakhsh, N., Callan, R., Alam, M., Prvulovic, M., and Zajic, A., "Leveraging electromagnetic emanations for iot security," (05 2017). Hardware Demo at IEEE International Symposium on Hardware Oriented Security and Trust (HOST).

[33] Alam, M., Khan, H. A., Dey, M., Sinha, N., Callan, R., Zajic, A., and Prvulovic, M., "One&done: A single-decryption em-based attack on openssl's constant-time blinded rsa," in [*27th Usenix Security Symposium*], (08 2018).

Approved for public release; distribution is unlimited