Introduction to the FIRST Special Issue

ANDREW CORMACK, Chief Regulatory Adviser, Jisc, Didcot, UK

JEROEN VAN DER HAM, Assistant Professor, Design and Analysis of Communications Systems Group, University of Twente, Enschede, The Netherlands, National Cyber Security Centre, The Netherlands

The Forum of Incident Response and Security Teams (FIRST) has held annual conferences around the world since 1989. They have always been fruitful places to exchange ideas: many developments in security and incident response can be traced back to discussions at FIRST conferences. For the first time, the 2019 conference in Edinburgh specifically targeted academic researchers, hoping to broaden these discussions further. We wanted to enable incident response practitioners to learn from the latest research and researchers to test their ideas against the practical experience of more than 1,000 conference delegates.

Presenters were invited to submit journal papers describing their work for this special issue of *ACM Digital Threats: Research and Practice*. The range of topics—from hardware exploits to process improvements—gives some impression of the breadth of the conference. We are particularly pleased that the authors include both academics and practitioners: a sign, perhaps, that our goal of enabling cross-fertilisation was successful.

Papers describing attacks on cryptographic systems can sometimes appear rather theoretical. In “[SCA-Pittaya: A Practical and Affordable Side-Channel Attack Setup for Power Leakage-Based Evaluations](#)”, François and Marc Durvaux describe a very practical experiment to recover cryptographic keys from two common Internet of Things (IoT) devices. They first monitor the power consumption of the 8-bit and 32-bit Arduino boards, recording the electromagnetic radiation emitted as they perform a small number of cryptographic operations using the 256-bit Advanced Encryption Standard algorithm. The resulting signal traces are then subject to statistical analysis to recover the secret keys. Even using the simplest correlation power analysis technique, the key can be recovered with high probability from just 350 traces: with preprocessing, this can be reduced to approximately 50, corresponding to less than a minute of recording. Perhaps the most striking achievement is that this does

CCS Concepts: • **Security and privacy** → **Cryptanalysis and other attacks**; **Embedded systems security**; **Malware and its mitigation**; **Intrusion detection systems**; *Usability in security and privacy*; • **Social and professional topics** → *System management*;

Additional Key Words and Phrases: Incident response, threat intelligence

ACM Reference format:

Andrew Cormack and Jeroen van der Ham. 2020. Introduction to the FIRST Special Issue. *Digit. Threat.: Res. Pract.* 1, 1, Article 2 (March 2020), 3 pages.

<https://doi.org/10.1145/3380746>

Authors' addresses: A. Cormack, Chief Regulatory Adviser, Jisc, Didcot, UK; email: andrew.cormack@jisc.ac.uk; J. van der Ham, Assistant Professor, Design and Analysis of Communications Systems Group, University of Twente, Enschede, The Netherlands, National Cyber Security Centre, The Netherlands; email: j.vanderham@utwente.nl.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2020 Copyright held by the owner/author(s).

2576-5337/2020/03-ART2

<https://doi.org/10.1145/3380746>

not require specialist equipment: the radiation signal is recorded using a commodity device (costing less than €400) and a homemade probe in contact with the microcontroller chip; the analysis is performed in a few seconds on a standard laptop. This definitely qualifies as theory leaving the laboratory and moving into the workplace, factory or home.

Olivier Van der Toorn and Anna Sperotto, in “[Looking Beyond the Horizon: Thoughts on Proactive Detection of Threats](#)”, explore whether the increased complexity of networked attacks may provide new opportunities for defenders. Several kinds of common malicious activity require attackers to assemble their own infrastructures before beginning the actual attack. The nature of these infrastructures means that they have to be publicly visible in the Domain Name System (DNS). If defenders could reliably identify these preparatory steps, they might be able to take proactive measures, before attacks are launched, to reduce their impact. Using the OpenINTEL project’s daily scans of well over half of the global DNS space, the authors investigate whether such identification is possible for three different types of attack and, if so, whether it could be used to respond earlier than current methods. For snowshoe spam, a comparison with current blacklists suggests promising results: 93% of the detected domains appear on current blacklists, suggesting good accuracy. The remainder may be either false positives or genuine attack infrastructures that were not used and blacklisted within the measurement period. More than 5% of these domains are detected from DNS data at least 2 days before they appear on a blacklist, and for 3% the DNS-based warning comes at least a week earlier. DNS domains may also be created for the purpose of launching traffic amplification attacks: here, a case study suggests that there is a distinct pattern created by the setting up of such a domain, but more work is needed to determine the parameters that could be used to automatically detect such patterns and how much prior warning this would provide. Finally, for combo-squatting attacks, the researchers conclude that although they can detect such domains, their typical lifecycle of rapid registration, use and disposal means that a daily scan of DNS records does not give significantly earlier warning than current approaches. Nonetheless, detecting attack infrastructures under construction seems a promising addition to the defender toolkit.

In “[Machine-Learning Framework to Analyze IoT Malware Using ELF and Opcode Features](#)”, Chin-Wei Tien, Shang-Wen Chen, Tao Ban, and Sy-Yen Kuo also explore patterns that may distinguish between legitimate and malicious activity, but in this case the source is executable code for IoT devices. Although static analysis of executable images is relatively common, previous tools have had to treat each instruction set separately. This leads to very large sets of patterns to be tested against, particularly for malware targeting IoT devices, where there is a particularly wide variety of different platforms and instruction sets. To produce a method that can be applied across multiple instruction sets, machine instructions are first divided into 12 categories that are common to all architectures, such as Logic, Control and Status, Memory and Stack. The frequency with which each of these categories appears in the binary plus seven other features of the Linux ELF file, such as size, architecture and whether it appears to use network communications, provide sufficient features to train a machine learning algorithm that is independent of architecture. The resulting system is able to classify previously unseen executables captured from IoT honeypots as either malware or benign with an accuracy of 97%, irrespective of the instruction set they use. The authors also experiment with labelling their training sets to see whether the model can assign malware to one of 15 families identified by VirusTotal (Gafgyt, Mirai, BitCoinMiner, etc.). This classifier also gives high accuracy. However, the available training and testing sets are not yet sufficiently diverse to conclude whether this is actually using features to identify the malware family across architectures or, where feature sets are similar between malware families, if the system may simply be selecting the family most commonly seen on the architecture of the code being inspected. Yet preliminary results suggest that with a sufficiently diverse training set, a satisfactory classification performance will be achieved.

In “[ARIMA Supplemented Metrics for Quality Assurance and Situational Awareness](#)”, Jan Kohlrausch and Eugene A. Brin observe that many parameters of interest to incident responders take the form of a series of measurements over time. These include signs of attacks (e.g., rate of scanning), the organisation’s own behaviour (e.g., time to patch) and issues with suppliers (e.g., quality of threat intelligence feeds). They propose that a

popular statistical method for analysing time series—autoregressive integrated moving average (ARIMA)—might be used to extract useful knowledge, such as means, trends and anomalies, from all of these data sources. The first requirement is to confirm that measurements have the characteristics required to make a time series suitable for ARIMA: they should be taken at regular intervals, sufficiently often for prompt detection of significant features, but not so frequently that they contain many zeros; additionally, a sufficient length of representative record (at least 50 points) should be available for initial training and periodic retraining. Because ARIMA is a well-established, well-analysed technique in other fields, statistical tests exist to check that data sources meet these requirements: the authors confirm that these are satisfied for point measurements of threat intelligence feeds and DSIELD scanning data, and suggest that they should also be met for measures such as mean time to respond when put through ARIMA’s differencing transformation. Experiments confirm that the technique can detect sudden changes in the volume of data arriving at a threat intelligence platform, and that it appears to set meaningful dynamic bounds for anomaly detection in port scan data. The latter measure was not positively tested due to the absence, confirmed by other sources, of any new worm activity targeting the chosen port during the experimental period. Demonstrating that typical datasets of interest to computer security incident response teams are suitable for ARIMA analysis suggests that adopting it as a general approach might have benefits ranging from financial and resource planning to automated threat detection.

Finally, in “[Fingerpointing False Positives: How to Better Integrate Continuous Improvement into Security Monitoring](#)”, Desiree Sacher suggests that when organisations dismiss security alerts as false positives, they are ignoring a rich source of information about their own security operations. Replacing the normal binary description—true/false positive—with eight different categories produces a classification that is meaningful across the organisation rather than local to each team’s interests. One of the new categories highlights where preventive security measures are working as intended and another where they are not. The other six indicate opportunities to improve the organisation’s record keeping, internal communications, configurations or procurement choices. Finding such valuable information in security teams’ least popular records could be a significant discovery.

As associate editors for this special issue, we are particularly pleased that all of these works suggest avenues for future development in both the theory and practice of incident response, whether in the need to provide better protection of IoT encryption, new sources of early attack detection or new opportunities to extract actionable information from undervalued datasets. It is our hope that theoreticians and practitioners will continue to work together on these and other ideas arising out of the 2019 FIRST Conference.

We wish to thank the authors, reviewers and programme committee members who have contributed to our text, the ACM for offering a special issue to host it and FIRST for sponsoring the issue so that it can be made open access.