Review of a Doctoral Thesis at FIT BUT

Doctoral thesis (hereinafter referred to as "thesis"), title of the thesis:

SECURITY OF WIRELESS COMMUNICATION FOR IOT DEVICES

Name of the doctoral student (hereinafter referred to as "candidate"), name and surname:

Ing. ONDŘEJ HUJŇÁK

Name and institution of the reviewer (full name of the reviewer, full name and country of the institution):

doc. Petr Švenda, Ph.D.

Masaryk University, Faculty of Informatics

Czechia

I. Thesis

Appropriateness and relevance

The thesis focuses on the domain of Internet of Things (IoT) security, which is very relevant for several reasons. The increasing number of such devices used, different approaches tried by vendors, and resource limitations posed by these platforms prevent the direct translation of well-established security mechanisms from the standard internet ecosystem. As a result, both analysis of current weaknesses and creative defensive methods are needed, including rigorous evaluation.

A summary of the contributions of the thesis

The thesis consists of eight primary chapters, focusing on resource-limited IoT devices for home and SOHO automation and covering a range of topics centered around the security of radio communication stack and the possibility of low-cost monitoring of ongoing network attacks. The assumptions about the difficulty in systematically fixing vulnerabilities found in loosely connected IoT devices are sensible and motivate the focus on monitoring ongoing attacks well.

Chapters two to four primarily map the existing state of the art and survey existing attacks and defenses, with chapters five and seven providing the candidate's own research contributions with clearly formulated research questions. Chapter six describes sound cryptographic design utilizing known security primitives to improve existing communication protocols but without clear research contribution. I will, therefore, provide comments primarily for the chapters with explicit research questions.

Chapter two provides a survey of wireless standards, their network parameters and configurations, and a detailed listing of relevant security and cryptographic features. Chapter three maps security features in IoT networks, covering generic attack taxonomies and focusing in detail on known attacks against LoRaWAN, ZigBee, and Bluetooth protocols. The mapping and survey of state of the art is very well written,

providing good insight into the current state of security of these protocols and their implementations, especially for Bluetooth.

Chapter four focuses on possible defensive measures, including monitoring, which is the primary focus of this thesis. It is argued that currently available IoT platforms focused on data processing are insufficient for more fine-grained monitoring due to their limited visibility of in-network traffic. Machine learning methods typically utilized in cloud or resource-abundant environments are also difficult to apply in-network due to the severely limited computation resources of IoT devices. Classification of network intrusion detection systems is provided. As a result, an alternative intrusion detection system is proposed and implemented based on the NEMEA framework.

Chapter five provides the core of the thesis with a focus on a very low-cost Bluetooth LE security monitoring platform. The main obstacle of external BT monitoring is the utilization of frequency hopping, which can be overcome by monitoring the whole ISM spectrum or more selective monitoring. The proposed monitoring technique observes BLE Advertising channel(s) for the periodic broadcasts of unconnected devices and then infers the time of connection based on the change of activity in the advertising channel.

Question for defense: The detection setup was tested in a controlled environment with three devices placed in a triangular shape roughly 50 centimeters away. How would the detection capability change with increased distance from the monitored BT device, as 50 centimeters is likely too close a distance for real-world deployment?

Q: The measurements based on ESP32 modules were compared to measurements done by the RPi BLE collector, with only a 2% difference observed. This observation is used to make a conclusion about the real advertising profiles of the devices that were discovered. Is such a conclusion warranted, or are other alternative explanations possible? For example, both ESP32 and RPi BLE modules may be based on similar enough underlying hardware, which would result in mutually small measurement differences, yet possibly large differences from true advertising profiles. The usage of high-end spectrum analyzers (even for a limited time, or possibly even a cheap SDR like RTL-SDR) as a source of ground truth would be beneficial.

Later, in section 5.7, more advanced outlier detection methods for connection detection of devices exhibiting intermittent advertising patterns are investigated. A grid search is performed to identify well-performing hyperparameters for different machine-learning approaches. The resulting comparison is extensive but, unfortunately, has not yet been scrutinized by an independent review process under multiple reviewers – the corresponding research publication is currently only in the submission stage. Also, given the machine learning methods used, I would suggest refraining from the usage of the overused term ‘artificial intelligence. ‘

Chapter six describes the candidate’s contribution to the design of security features of the IQRF wireless platform. The previous (rather weak) cryptographic protections are described (IQRF Legacy), followed by a detailed description of a significantly improved version utilizing well-known cryptographic primitives (AES, AEAD) and proactive key management techniques. The chapter is reasonably well written; however, the research question is not explicitly stated, and its academic contribution is not clear.

Chapter Seven investigates several research questions in the domain of privacy leaks observable from communication between IoT gateways and public networks. Traffic from 4 different gateways was captured for a period of 28 days and analyzed. The interesting observations concluded that commercial

devices generate more telemetry traffic than open-source ones. The hypothesis about the potential scanning of UPnP devices was formulated, but it was not investigated further due to traffic being encrypted. Would it be possible to test for such behavior by controlled addition and removal of UPnP-enabled devices into the internal network, followed by observation of changes in encrypted communication metadata (especially data length)? This observatory analysis is followed by a discussion of several options for device anonymization. The generic anonymization approaches are described with a strong focus on the Tor mixing network. The additional value of the analysis is in the application to the specifics of the IoT networks. The security analysis is performed, but it is rather informal, with claims like “”…this seems to be secure.” (pp 115) not being rigorous enough. The second part of the chapter (7.2 Anonymization Strategies) provides a good starting point but deserves a more thorough analysis.

I value the practical execution of numerous experiments to obtain actual, real-world data for subsequent analyses. The source codes for a proof of concept devised in Chapter 5 are publicly available (although the license does not seem to be explicitly specified). The corresponding dataset from the monitoring of the BLE advertisements channel was also published. The code for the evaluation of machine learning techniques comparison seems not to be available yet, but I suppose this is due to the corresponding paper still being under review - consider adding an explicit note that the code will be publicly released if that is the case.

Overall, I believe that the thesis achieved its goal, especially when the research currently under review would be accepted.

Novelty and significance:

The candidate’s research contribution is relevant and provides a deeper understanding of the usability of Bluetooth advertisements for low-cost detection of Bluetooth device usage periods. Although the general technique is already utilized by special commercial systems intended for monitoring of BT, WiFi, or mobile network transmissions and used at airports, shopping malls, and public buildings, mostly for movement monitoring and commercial marketing purposes (e.g., BlipTrack, Walkbase), the exact details are not public and rigorous evaluation was missing – the gap filled by this thesis and papers under submission. Similarly, the monitoring of potential privacy leaks on IoT gateways is not a novel technique itself, but the work done in this thesis provides more insight, discovers previously unknown information, and discusses potential mitigations.

Evaluation of the formal aspects of the thesis:

The thesis is written in the English language and is understandable with no significant problems, with only some smaller stylistic issues (interchangeable usage of cypher vs. cipher, widow (pp.30), duplicate reference for NEMEA framework [31] and [126], third-person form in chapter 4.4). The text is (understandably) compiled from the candidate’s own papers with subsequent improvements and reorganizations with some minor deficiencies – for example, page 13 talks about the paper instead of the thesis. The text is logically structured, and enough explanatory details are provided.

Quality of publications

The core contributions of the thesis were published in nine conference and journal articles. Out of these, three were at an international conferences with Core B rank (103% contribution total), one at WoS Q2 journal (30% contribution), and the remaining five at unranked, frequent country-local conferences. The

papers tend to be of a shorter page format. Additional two papers were still under submission at the time of this review writing.

Overall, the quantity of the candidate's research output publications is sufficient, especially if two papers currently under submission are accepted. The quality of the venues selected for publications is moderate, but with positively improving tendency. The initial papers were published at unranked venues and shorter paper formats, but more recent papers in 2023 were at Core B conferences. Given the first paper was published already in year 2017, I would suggest to consider more ambitious venues in an earlier stage of his postgraduate study. As the candidate demonstrated the ability to carry quality research in later stages, earlier focus on more demanding venues would possibly result in faster progress, better feedback from reviewers and improved outputs overall.

II. Student's overall achievements

Overall R&D activities evaluation:

The candidate demonstrated his ability to conduct original scientific research, formulate the research question clearly, design experiments, and evaluate the results. The description of the methodology, developed tools, and collected data are described clearly and released as open-source, helping to make the research replicable. The combination of already published papers, together with the additional two papers currently in submission, show the ability of a candidate to produce research output of good quality.

Assessment of other characteristics (optional):

The candidate's expertise in security designs and known attacks in the IoT domain was also used to help design improved security specifications of IQRF protocol as described in Chapter 7. While such application of expertise itself does not fall directly into the category of original academic research, it demonstrates the candidate's ability to translate research into practice.

III. Conclusion

In my opinion, the thesis and the candidate´s achievements meet the generally accepted requirements for the award of a Ph.D. academic degree, especially when two papers currently under submission are considered but also without them. The candidate's results achieved clearly improved as his study progressed, resulting in submissions to more competitive publication venues. I value the candidate's attitude towards replicable research by making source code and datasets publicly available.

Brno 31.10.2024

Signature of the reviewer: