Towards In-Network Security for Smart Homes

Martin Serror*, Martin Henze*, Sacha Hack§, Marko Schuba§, Klaus Wehrle*
*Communication and Distributed Systems, RWTH Aachen University, Germany
§FH Aachen University of Applied Sciences, Germany
{serror,henze,wehrle}@comsys.rwth-aachen.de,{s.hack,schuba}@fh-aachen.de

ABSTRACT

The proliferation of the Internet of Things (IoT) in the context of smart homes entails new security risks threatening the privacy and safety of end users. In this paper, we explore the design space of *in-network* security for smart home networks, which automatically complements existing security mechanisms with a rule-based approach, i. e., every IoT device provides a specification of the required communication to fulfill the desired services. In our approach, the home router as the central network component then enforces these communication rules with traffic filtering and anomaly detection to dynamically react to threats. We show that in-network security can be easily integrated into smart home networks based on existing approaches and thus provides additional protection for heterogeneous IoT devices and protocols. Furthermore, in-network security relieves users of difficult home network configurations, since it automatically adapts to the connected devices and services.

CCS CONCEPTS

• **Security and privacy → Network security**; *Intrusion/anomaly detection and malware mitigation*;

KEYWORDS

Internet of Things, Smart Homes, Network Security, Attack Mitigation, Anomaly Detection

ACM Reference Format:
Martin Serror, Martin Henze, Sacha Hack, Marko Schuba, Klaus Wehrle. 2018. Towards In-Network Security for Smart Homes. In *ARES 2018: International Conference on Availability, Reliability and Security, August 27–30, 2018, Hamburg, Germany*. ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3230833.3232802

ARES 2018, August 27–30, 2018, Hamburg, Germany

ACM ISBN 978-1-4503-6448-5/18/08…$15.00
https://doi.org/10.1145/3230833.3232802

1 INTRODUCTION

With the proliferation of the Internet of Things (IoT) [2], all sorts of Internet-connected devices, ranging from simple sensors to more complex controllers and home appliances, find their way into private homes [14]. For the convenience of their users, these IoT devices are either directly or indirectly (via a bridge) connected to the home network enabling continuous monitoring and control through the Internet. This trend, referred to as *smart homes*, encompasses home appliances, smartphones, heating and cooling, and alarm systems [33]. As a result, users benefit from a streamlined automation of daily tasks, enhanced monitoring and alarm capabilities, and even reduced energy consumption. These benefits have led to the unabated success of smart home deployments, e. g., Gartner predicts that 20 billion devices will be connected to the Internet by 2020, where consumer devices will represent the largest group with around 13 billion anticipated devices [20].

Besides these benefits, however, smart homes introduce serious security challenges [6, 12, 15, 24], which mainly result from the unrestricted interconnection of IoT devices among each other and with the Internet [16]. Since the local communication of IoT devices is often unencrypted [7], devices rely on well-known standard passwords [31], or enable remote access, e. g., via `telnet`, an attacker can exploit these security flaws to gain unauthorized (root) access to devices [29]. Once the attacker gained access, further attacks are possible, such as eavesdropping on network traffic to access sensitive information, manipulating devices, e. g., deactivating an alarm system, or even installing ransomware [43]. As a result, a single compromised device in a home network, e. g., a malware-laden smartphone [36], often suffices to gain access to other devices in this network. Considering currently deployed IoT devices, an Internet-wide analysis of reachable IoT devices shows that the vulnerability rate strongly depends on the device type, varying between 0.44 % and 40 % [31]. Recent attacks, such as Mirai [1], showed that security flaws of IoT devices have far-reaching consequences, since these devices can be turned into a powerful botnet performing Distributed Denial-of-Service (DDoS) attacks on vital online services.

We observe that these security risks are often exacerbated by the methodology of underlying home network deployments where, once a device is connected to the home network, there are no mechanisms on the network layer that restrict the communication capabilities of this device to the bare minimum required for the intended functionality. For example, if a smart home device such as a web radio gets infected by malicious software, it can establish new connections and attack other devices and servers without limitations. However, conceptually there is no need to allow communication of a web radio with other devices to realize the intended functionality, i.e., streaming audio from a distinct number of Internet hosts. From a different perspective, typical users of IoT devices are often overwhelmed with securely configuring their smart home network. The new possibilities enabled by a large variety of IoT devices and services invites users to connect devices with often lenient default security configuration to their home network, without thinking about the consequences, or simply because they do not know about the potentially involved risks.

In this paper, we address these security risks that are imminent to the current deployment model of smart homes and IoT devices in home networks. To this end, we argue that security in smart homes must be addressed *automatically* and within the local network in addition to the security mechanisms already deployed directly at the device or service side. Notably, our proposed approach of *in-network* security offers orthogonal protection to the security mechanisms implemented on the devices by monitoring the local network, detecting suspicious traffic, and preventing attacks from and to IoT devices connected to the home network.

In contrast to the state-the-art in this area [4, 8, 18, 34, 35, 37], we believe that the communication of IoT devices must be restricted to the bare minimum required for the intended functionality to *proactively* prevent security breaches. To this end, we propose to equip each IoT device with a specification of the required communication to fulfill its services. Consequently, all other traffic is blocked and a device's communication is monitored for non specification-compliant behavior. Our approach thus differs from existing network security approaches such as intrusion detection [28], as it builds upon explicit permissions to prevent unauthorized network traffic instead of generously allowing all communication first and then trying to detect anomalous behavior.

In particular, our contributions lie in a thorough description of the security threats in current smart homes as well as a clear motivation for in-network security in such deployments (Sec. 2). Based on our analysis, we propose three main components for realizing in-network security:

(1) Specification of network compliant behavior for individual IoT devices in smart home networks (Sec. 3);
(2) IoT traffic filtering to enforce compliance with the specified communication behavior (Sec. 4); and
(3) anomaly detection to be able to dynamically counter attacks within the smart home network (Sec. 5).

In the following, we explore the design space for each of these components by focusing on existing as well as promising future approaches to realize individual parts of in-network security, before we point out future research challenges (Sec. 6).

2 PROBLEM ANALYSIS

The main challenges for securing smart home networks arise from the heterogeneity of IoT devices regarding their capabilities, modes of operation, and applications [19]. With an increasing number of IoT devices, users are overburdened with configuring and maintaining each device in a secure manner [17]. Typically, it is often unclear what kind of network permissions a specific device should have. As a result, users tend to blindly trust the IoT devices they connect to their home network. In the following, we illustrate the potential security threats in smart homes and derive the necessity for an in-network solution to mitigate these threats.

2.1 Security Threats in Smart Homes

Smart homes introduce serious security threats ranging from exposing sensitive information to unauthorized third parties [15] over hacked devices that are used to launch DDoS attacks against innocent bystanders [1, 3] to physical harm, e. g., by manipulating a smart lock to break into a home [11]. In response, the research community invested enormous efforts to identify, describe, and categorize these security threats [6, 8, 21, 24, 25]. We provide an overview of the most important security threats in smart homes, especially with respect to network communication, alongside an example scenario in the following.

In our (fictitious) example, a user buys an IP camera to monitor the drive way in front of her home. To quickly configure the camera, she uses her smartphone to connect to the camera via Bluetooth. Once the connection is established, the configuration menu of the camera allows the user to select her home network from a list of nearby WiFi networks and subsequently to enter the corresponding password. The IP camera is now connected to her home network and has full access to the local network and to the Internet, just like any other connected IoT device. To allow the user to access her camera remotely, the camera automatically and unknowingly to the user opens a port in the firewall of the user's router, e. g., using Universal Plug and Play (UPnP).

It is, however, unknown to the user that the camera is globally reachable and that the current firmware image includes several security flaws, e. g., well-known standard passwords [31] or the potential for remote code execution which

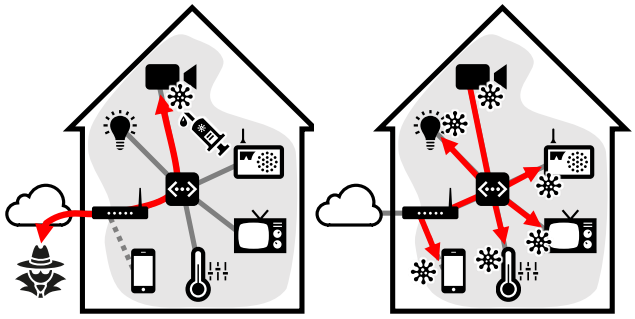

(a) An attacker gains access due to security flaws in the IP camera.

(b) The infected IP camera attacks other devices in the home network.

Figure 1: Attack scenario in which an infected IoT device attacks other devices within the same network.

can be exploited to gain unauthorized (root) access to the camera [41]. These security flaws enable an attacker to gain access to the IP camera from outside, and even to install malware on the camera. In our example, depicted in Fig. 1, the malware then takes advantage from the fact that many other devices are connected to the same home network without any restrictions w.r.t. their permitted communication behavior. The malware on the camera thus scans the local network for other, potentially vulnerable devices and intends to get access to them, e. g., by exploiting other well-known security flaws, to install further copies of the same or a different malware. Once the infected devices are under the control of the malware, they may not only cause harm in the smart home network itself, e.g., by exposing private information, but also perform other malicious activities such as contributing to DDoS attacks [3] or manipulating services.

Although this is only a single example of a possible attack vector against a smart home, it, nonetheless, illustrates the fundamental weaknesses of current smart home network configurations, which can be exploited in numerous ways [12]. Notably, to perform such attacks, it is not even necessary that a vulnerable device is directly reachable from the Internet (as in our above example). Recent research shows, e.g., that a malware-laden smartphone can be exploited to attack devices in an otherwise secured home network [36]. As a result, typical protection mechanisms at the network layer, e. g., a firewall, that aim at blocking unauthorized network traffic from the outside are no longer sufficient to protect devices in smart home deployments. As soon as one device in a home network is under the control of malicious software, there are nowadays practically no security measures to prevent attacks towards other devices and services originating from the infected device.

2.2 The Need for In-Network Security

To augment security in smart homes, we thus require an *in-network* approach that automatically adapts to the heterogeneity of smart home networks by restricting the communication capabilities of IoT devices without limiting their desired functionality. To this end, we propose a new approach to network security that restricts both internal communication (i.e., with other devices in the same home network) *and* external communication (i.e., with Internet- and cloud-based services) of individual IoT devices to the extent necessary for delivering their intended functionality. That is, depending on the purpose of the IoT device in the smart home network, we allow certain incoming and outgoing connections of the device that are necessary for its intended functionality, while blocking all other connections.

Introducing this new network security approach serves several purposes. First of all, the unauthorized access to IoT devices that are connected to the Internet is prevented by blocking such connection attempts. Moreover, in the case that an attacker still gained access to an IoT device, e. g., by exploiting a security flaw, the unauthorized access to other devices within the network should be impeded. Existing network security mechanisms, e. g., a firewall, typically operate on the home router protecting the network from unauthorized access from outside. In turn, for connections within the home network such security mechanisms do not take effect, which allows infected devices to easily attack other devices within the network. Therefore, our proposed communication approach only permits those connections within the network that are necessary such that the IoT devices can fulfill their tasks. Hence, we significantly reduce the attack vector of rogue devices *within* the home network. Finally, unauthorized connections from the home network to other networks are also blocked, e. g., to prevent infected devices to perform a DDoS attack or expose sensitive information that should not leave the home network.

To turn our vision of a new communication approach for in-network security of smart homes into reality, we require three main components as illustrated in Fig. 2: ① a *specification of network compliant behavior* for each IoT device to derive what kind of communication has to be performed to deliver the desired functionality, ② a centralized network component, e. g., the home router, that *filters IoT traffic* according to the specified intended communication behavior to prevent attacks from and to IoT devices, and ③ an *anomaly detection system* that enables us to dynamically react to attacks within the smart home network by evolving the specified intended communication behavior.

Our approach is based on the assumption that all communication traffic is routed through a central network router, which is the typical setting in today's home networks. This,

Service Type	Description	Source	Destination	Traffic Type	Max. Packet Size	Max. Data Rate
audio stream	requests & ACKs	`*.*.*.*:*`	`example.org:80`	periodic	200 Bytes	8 kbit/s
	AAC data	`example.org:80`	`*.*.*.*:*`	periodic	*	128 kbit/s
station info	requests & ACKs	`*.*.*.*:*`	`example.org:80`	bursty	200 Bytes	*
	XML data	`example.org:80`	`*.*.*.*:*`	bursty	*	*
firmware update	requests & ACKs	`*.*.*.*:*`	`203.0.113.10:21`	periodic	200 Bytes	8 kbit/s
	binary data	`203.0.113.10:21`	`*.*.*.*:*`	bulk	*	*
remote config	https	`cloud.example.org:*`	`*.*.*.*:443`	bursty	600 Bytes	*
	https	`*.*.*.*:443`	`cloud.example.org:*`	bursty	600 Bytes	*

Table 1: Example of a service description for a web radio.

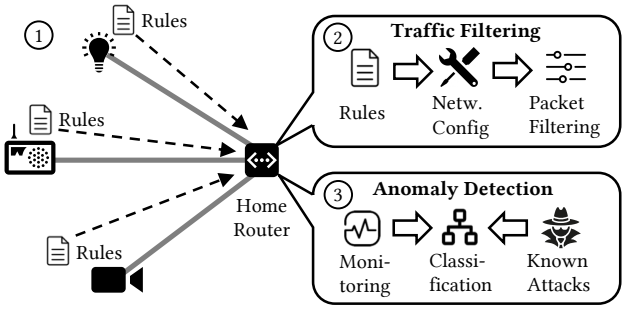

Figure 2: Overview of in-network security: ① The communication rules of each device are installed at the home router. ② These rules are used to filter malicious traffic. ③ Moreover, anomaly detection is used to detect new attack vectors.

however, implies that attacks occurring in separated networks, e. g., via Bluetooth, can only be detected by our approach when they also affect the network managed by the home router. Furthermore we assume that the soft- and hardware of the home router has not been compromised and that (external) attackers perform large scale attacks instead of targeting specific users or networks.

In the remainder of this paper, we detail these three components and discuss how they can be realized based on prior work of the research community as well as promising future approaches.

3 SPECIFICATION OF NETWORK COMPLIANT BEHAVIOR

An essential design component of our envisioned smart home in-network security approach is the specification of network compliant behavior for each IoT device to distinguish malicious from harmless, i.e., intended, network traffic. As stated before, compliant behavior of an IoT device entails the minimum required communication to fulfill the desired services. We depict an example for the provided services and required communication of a web radio in Table 1.

To leverage knowledge of network compliant behavior as a foundation to realize in-network security for smart homes, we propose to specify this behavior in the form of *communication rules*. Communication rules specify the IP addresses (or hostnames) as well as port numbers and communication direction that need to be permitted to realize the intended functionality of a specific IoT device, similar to, e.g., rules for firewalls. Furthermore and in contrast to such simple firewall rules, communication rules also define other essential characteristics of network communication such as packet sizes, packet interarrival times, number of parallel connections, and consumed bandwidth. Each device can have multiple communication rules to account for different intended functionality of a single device, e. g., a web radio (cf. Table 1) typically receives audio streams at a bandwidth of 128 kbit/s but occasionally also needs to download new station lists or firmware updates. Moreover, the user may configure her web radio via a cloud service that interacts with the device via `https`.

In its easiest form, such communication rules specifying network compliant behavior are directly provided by the manufacturer of the respective device. To account for legacy devices already deployed in smart homes today or manufacturers that refrain from publishing communication rules for their devices, users can rely on a certification agency or the open-source community to provide communication rules, or even provide them by themselves [42]. To this end, recent research shows that IoT devices can be automatically identified by passively observing network traffic [27]. Such fingerprinting techniques can be used to retrieve the respective communication rules from a trusted third party, e. g., a certification agency or an open-source repository, given that the manufacturer does not provide a description of network compliant behavior in the first place.

Besides identifying individual IoT devices, we can also classify devices into groups based on their communication behavior. As a result, we can provide certain default rules to a whole group of devices with similar communication behavior. For example, when a new IoT device, for which no communication rules are available (not even from a trusted

third party), connects to the home network, this device can be automatically matched to an existing communication class, e. g., *permit only local communication* or *permit only communication to cloud services*. This enables more specified default rules for unknown devices, which can be gradually adapted over time when more traffic data becomes available.

Especially if the communication behavior of a new device does not closely match one of the default classes, it might be necessary to automatically derive the communication rules of this device. To this end, we propose to observe the communication of the new device during a learning phase and create communication rules based on the initial behavior of this device. Here, it is important to cover the different aspects of the device's functionality (see above) to reduce the overhead for adapting communication rules later on.

Once we have obtained the communication rules for an IoT device through one of the above channels, these rules are centrally installed at the home router that bridges the home network with the Internet. In the case that such a description is not available yet, e. g., because we first need to classify the IoT device based on its communication behavior, we apply a default set of rules. As such, the device is, e. g., not allowed to open a connection to another device or server unless the user explicitly grants permission at the home router.

For a practical realization of such a system, we expect it to be challenging to implement the real-time traffic filtering on commercial off-the-shelf router hardware. Nevertheless, as shown in the following section, there are existing approaches that target efficient traffic filtering. For future work, we thus aim at adapting these approaches to fit the requirements and limitations of the smart home use-case.

4 IOT TRAFFIC FILTERING

Based on the communication rules that specify network compliant behavior for each IoT device, we propose proactive *in-network* security that aims at minimizing the risk of attacks within the smart home network. This not only includes attacks from outside but also attacks from inside, e. g., by infected IoT devices connected to the smart home network. In the following, we describe how we envision the realization of in-network security based on a Software-Defined Networking (SDN) and flexible packet matching approach that efficiently enforces the communication rules in a heterogeneous environment with evolving attack vectors.

4.1 SDN-Based Attack Prevention

Software-Defined Networking (SDN) realizes central network management, which separates the data plane from the control plane [23]. This facilitates the change of network protocols and policies in existing networks, since only the software of the central controller needs to be adapted. Especially large networks, e.g., in data centers and companies, benefit from SDN due to the reduced complexity and overhead for deployment and maintenance. However, the advantages of SDN have been also considered for automating the security of home networks [10] and thus to relieve users of complex security management tasks.

In alignment with previous research efforts [34, 37] we therefore propose to follow an SDN approach in smart home networks to enforce network compliant behavior as defined by our communication rules (cf. Sec. 3). This approach does not require additional hardware, since the SDN controller can be centrally located at the (programmable) home router. As a first step for IoT traffic filtering, the communication rules for each IoT device need to be converted to flow-table entries for the forwarding tables, specifying for each data flow how it is matched and which actions have to be taken. If, for example, the communication rules for a smart lighting system specify only local communication within the home network, then all packets originating from the lighting system with a destination IP that does not lie within the home network address range are dropped. Besides dropping packets, other actions according to the communication rules are possible, e. g., limiting the traffic rate of an IoT device.

The flexibility of SDN allows to dynamically adapt the communication rules, e. g., when new devices are connected with the home network or the functionality of existing devices is extended. Even further, the enforcement of such rules can be realized on the IoT devices themselves and on the smartphones that interact with these devices [8]. Thus, we can filter malicious network traffic either directly at its source or at least close to its origin. However, the main challenge to realize this SDN-based approach is to correctly filter network packets according to the communication rules of individual devices, which we further elaborate in the following.

4.2 Flexible Packet Matching

The de facto standard protocol for configuring SDN deployments is OpenFlow [38]. However, its inflexible packet matching based on hard-coded header fields makes it unsuitable for quickly evolving protocols [5]. For smart home networks with a broad range of (unknown) IoT devices and protocols, the traditional OpenFlow deployment model needs to be evolved to enable dynamic packet matching. In the following, we shortly present two distinct approaches that enable flexible packet matching in SDN and thus are well-suited for enforcing security rules in smart home networks.

P4. To enable flexible packet matching, Bosshart et al. [5] propose Programming Protocol-independent Packet Processors (P4), a high-level language that allows to specify the behavior of a network switch. More specifically, P4 can be used

to describe, independently of the target hardware, a parse-match-action pipeline, i. e., a description of packet matching sequences, packet processing, and actions, which can all be modified during run time. Therefore, the P4 compiler first translates the P4 program into multiple match+action tables, which are then analyzed against dependencies with the help of table dependency graphs to enable sequential and parallel execution. Afterwards, the compiler adapts this representation for the target soft- and hardware. P4 thus allows to specify match+action rules independently of the hardware and, more importantly, to flexibly adapt to evolving protocols and security threats. For in-network security in smart homes, each IoT device can thus specify its communication rules for network compliant behavior in a small P4 program including the required headers, parsers, tables, and actions. Subsequently, this program is used to filter the traffic of the IoT device to permit only network compliant behavior.

eBPF. Berkeley Packet Filters (BPF) were originally specified for efficient packet filtering in the Linux kernel already in 1993 [26]. Until today, BPF evolved to a powerful tool for network monitoring, engineering, and security; especially since its major revision known as extended BPF (eBPF) [13] has been included into the Linux kernel in 2014. More specifically, eBPF allows to specify small filter programs, which are executed in a virtual machine in the operating system kernel, close to the actual hardware. Since running user code in the kernel poses certain risks, eBPF includes a verifier that performs several security and stability checks before loading an eBPF program [32]. Due to its efficiency and powerfulness, Jouet et al. [22] propose to use eBPF for flexible packet matching in OpenFlow. Similarly to P4, eBPF allows to match packets independently of any protocol implementation and is therefore well-suited for the heterogeneous IoT landscape. Furthermore, Jouet et al. show that eBPF reduces the number of flow entries and achieves efficient packet filtering at line-rate. In contrast to P4, changes in the eBPF program specification do not entail a recompilation of the target switch software [39]. For our approach of in-network security for smart homes, eBPF thus allows to easily implement and modify communication rules of network compliant behavior during operation. Such updates become necessary when new IoT devices are connected with the home network, or in the case of anomalous behavior, e. g., when a device gets infected and intents to attack other devices. We detail our discussion on the need to detect and react to anomalies in the communication behavior in the next section.

5 ANOMALY DETECTION

So far, we described how our proposed approach for in-network security realizes communication rules for network compliant behavior (cf. Sec. 3) and enforces these rules in the smart home network based on the principle of SDN (cf. Sec. 4). However, taking into account the heterogeneity and the vast growth of the number of IoT devices, it is possible that new attack vectors evolve that were not considered in the existing set of communication rules. Furthermore, the provided communication rules might be incomplete and contain errors or still allow for exploiting unforeseen attacks. The proposed proactive security mechanisms of our approach therefore must be complemented with a reactive approach that dynamically adapts the current smart home configuration. To this end, we propose an efficient network monitoring system as well as anomaly detection based on machine learning.

5.1 Monitoring IoT Devices

A prerequisite for network-level anomaly detection is an efficient monitoring of the network traffic. In the context of smart home networks we expect, on the one hand, an increasing number of IoT devices and consequently increasing network traffic. On the other hand, typical networking hardware for domestic use, e. g., routers and switches, has limited computational resources. As a result, the monitoring approach can only focus on essential packet characteristics, such as packet headers, lengths, and periodicity, instead of performing deep packet inspection, i.e., analyzing the payload of individual network packets [18]. Still, such essential packet characteristics can be used to identify traffic patterns of IoT devices and thus to compare these patterns against known malicious behavior. To even further reduce the costs of monitoring in smart home networks, Sivanathan et al. [35] propose to monitor the network traffic at flow-level granularity, where not all packets of a data flow are analyzed in detail to keep the processing overhead low. Consequently, we propose to first analyze the initial packets of a new unknown data flow to quickly assess whether this flow is suspicious or not and to adapt the monitoring accordingly. Furthermore, we use the collected data post factum [40], e. g., to isolate a culprit device with the help of the recorded data once an attack has been detected.

5.2 Machine Learning Techniques

Besides detecting anomalies in the smart home network based on known traffic patterns, a more dynamic approach consists in using machine learning to classify data flows [30]. An advantage of using machine learning for anomaly detection is the possibility to detect unknown attacks, which were not considered in the system yet. A disadvantage is, however, the risk of misclassification leading to a too restrictive network configuration with reduced functionality. Therefore, a machine learning approach should only complement the aforementioned rule-based approach, e. g., by adapting existing communication rules or reacting on imminent attacks.

As proposed in [4], a learning module can be trained for each data flow during run time based on the monitored traffic (cf. Sec. 5.1). Additionally, known attack patterns, e. g., TCP SYN flooding attacks, can be fed into the learning module as training data [9]. A classification module then classifies data flows based on the input of the learning module, e. g., using decision trees, neural networks, random forests, or support vector machines [9]. In case an anomaly has been detected, the corresponding communication rules are adapted to prevent further risks to the smart home network.

6 CONCLUSION AND FUTURE RESEARCH CHALLENGES

In this paper, we propose automatic *in-network* security for smart homes to tackle the security challenges introduced by heterogeneous IoT devices and protocols, and to relieve users of complicated security configuration of their home networks. Therefore, we are convinced that each IoT device should describe the communication required to realize its functionality, such that these communication rules can be enforced in the smart home network to realize network compliant behavior. Based on existing SDN-based approaches and efficient packet matching engines, we explore the design space for realizing such a system in the context of smart homes. Furthermore, network compliance behavior-driven in-network security can be complemented with machine-learning techniques to dynamically react to imminent attacks and to adapt communication rules during operation.

For future work, we aim at a prototypical realization of in-network security for smart homes to empirically measure its performance and to show the security benefits of this approach. We are particularly interested in achieving a balanced definition of communication rules, which does not limit the functionality of the IoT device and at the same time achieves a strong level of security for the smart home network. In this context, the automatic classification of IoT devices into different groups based on their initial traffic is of special interest to support legacy devices, where no communication rules are available. Furthermore, we expect that integrating machine learning techniques into the rule-based system will significantly improve the detection rate of malicious network traffic, compared to systems that exclusively apply machine learning without any domain knowledge. Finally, we consider it promising to extend the concepts of in-network security to related but yet different environments, such as the Industrial Internet of Things (IIoT).

ACKNOWLEDGMENTS

This research was supported by the research training group "Human Centered System Security" sponsored by the German federal state of North Rhine-Westphalia.

REFERENCES

[1] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, and others. 2017. Understanding the Mirai Botnet. In *USENIX Security Symposium*. USENIX Association, Vancouver, BC, 1093–1110.

[2] L. Atzori, A. Iera, and G. Morabito. 2010. The Internet of Things: A Survey. *Computer Networks* 54, 15 (Oct. 2010), 2787 – 2805.

[3] E. Bertino and N. Islam. 2017. Botnets and Internet of Things Security. *Computer* 50, 2 (Feb. 2017), 76–79.

[4] S. S. Bhunia and M. Gurusamy. 2017. Dynamic Attack Detection and Mitigation in IoT using SDN. In *27th International Telecommunication Networks and Applications Conference (ITNAC)*.

[5] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker. 2014. P4: Programming Protocol-independent Packet Processors. *SIGCOMM Computer Communication Review* 44, 3 (July 2014), 87–95.

[6] J. Bugeja, A. Jacobsson, and P. Davidsson. 2016. On Privacy and Security Challenges in Smart Connected Homes. In *European Intelligence and Security Informatics Conference (EISIC)*. IEEE, 172–175.

[7] M. Capellupo, J. Liranzo, M. Z. A. Bhuiyan, T. Hayajneh, and G. Wang. 2017. Security and Attack Vector Analysis of IoT Devices. In *Security, Privacy, and Anonymity in Computation, Communication, and Storage*. Springer International Publishing, 593–606.

[8] S. Demetriou, N. Zhang, Y. Lee, X. Wang, C. A. Gunter, X. Zhou, and M. Grace. 2017. HanGuard: SDN-driven Protection of Smart Home WiFi Devices from Malicious Mobile Apps. In *Proceedings of the 10th Conference on Security and Privacy in Wireless and Mobile Networks (WiSec)*. ACM, 122–133.

[9] R. Doshi, N. Apthorpe, and N. Feamster. 2018. Machine Learning DDoS Detection for Consumer Internet of Things Devices. In *Deep Learning and Security Workshop (DLS)*. IEEE.

[10] N. Feamster. 2010. Outsourcing Home Network Security. In *Proceedings of the SIGCOMM Workshop on Home Networks (HomeNets '10)*. ACM, 37–42.

[11] E. Fernandes, J. Jung, and A. Prakash. 2016. Security Analysis of Emerging Smart Home Applications. In *IEEE Symposium on Security and Privacy (SP)*. 636–654.

[12] D. Geneiatakis, I. Kounelis, R. Neisse, I. Nai-Fovino, G. Steri, and G. Baldini. 2017. Security and Privacy Issues for an IoT based Smart Home. In *40th International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*. IEEE, 1292–1297.

[13] B. Gregg. 2015. eBPF: One Small Step. http://www.brendangregg.com/blog/2015-05-15/ebpf-one-small-step.html. (May 2015).

[14] M. Henze, L. Hermerschmidt, D. Kerpen, R. Häußling, B. Rumpe, and K. Wehrle. 2014. User-driven Privacy Enforcement for Cloud-based Services in the Internet of Things. In *Proceedings of the 2014 International Conference on Future Internet of Things and Cloud (FiCloud)*.

[15] M. Henze, L. Hermerschmidt, D. Kerpen, R. Häußling, B. Rumpe, and K. Wehrle. 2016. A Comprehensive Approach to Privacy in the Cloud-based Internet of Things. *Future Generation Computer Systems* 56 (2016).

[16] M. Henze, J. Hiller, R. Hummen, R. Matzutt, K. Wehrle, and J. H. Ziegeldorf. 2017. Network Security and Privacy for Cyber-Physical Systems. In *Security and Privacy in Cyber-Physical Systems: Foundations, Principles and Applications*, H. Song, G. A. Fink, and S. Jeschke (Eds.).

[17] M. Henze, B. Wolters, R. Matzutt, T. Zimmermann, and K. Wehrle. 2017. Distributed Configuration, Authorization and Management in the Cloud-based Internet of Things. In *IEEE Trustcom/BigDataSE/ICESS*.

[18] C. Hesselman, J. Jansen, M. Davids, and R. de O. Schmidt. 2017. *SPIN: a User-centric Security Extension for In-home Networks*. Technical Report SIDN-TR-2017-002. SIDN Labs.

[19] M. M. Hossain, M. Fotouhi, and R. Hasan. 2015. Towards an Analysis of Security Issues, Challenges, and Open Problems in the Internet of Things. In *IEEE World Congress on Services*. 21–28.

[20] Gartner Inc. 2017. Gartner Says 8.4 Billion Connected "Things" Will Be in Use in 2017, Up 31 Percent From 2016. https://www.gartner.com/newsroom/id/3598917. (2017).

[21] Y. J. Jia, Q. A. Chen, S. Wang, A. Rahmati, E. Fernandes, Z. M. Mao, and A. Prakash. 2017. ContexIoT: Towards Providing Contextual Integrity to Appified IoT Platforms. In *Proceedings of the Network and Distributed System Security Symposium (NDSS)*. Internet Society.

[22] S. Jouet, R. Cziva, and D. P. Pezaros. 2015. Arbitrary Packet Matching in OpenFlow. In *IEEE 16th International Conference on High Performance Switching and Routing (HPSR)*.

[23] H. Kim and N. Feamster. 2013. Improving Network Management with SDN. *IEEE Communication Magazine* 51, 2 (Feb. 2013), 114–119.

[24] N. Komninos, E. Philippou, and A. Pitsillides. 2014. Survey in Smart Grid and Smart Home Security: Issues, Challenges and Countermeasures. *IEEE Communications Surveys Tutorials* 16, 4 (April 2014), 1933–1954.

[25] R. Mahmoud, T. Yousuf, F. Aloul, and I. Zualkernan. 2015. IoT Security: Current Status, Challenges and Prospective Measures. In *10th International Conference for Internet Technology and Secured Transactions (ICITST)*. 336–341.

[26] S. McCanne and V. Jacobson. 1993. The BSD Packet Filter: A New Architecture for User-level Packet Capture. In *Proceedings of the USENIX Conference*.

[27] M. Miettinen, S. Marchal, I. Hafeez, N. Asokan, A. R. Sadeghi, and S. Tarkoma. 2017. IoT Sentinel: Automated device-type identification for security enforcement in IoT. In *IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2177–2184.

[28] B. Mukherjee, L. T. Heberlein, and K. N. Levitt. 1994. Network Intrusion Detection. *IEEE Network* 8, 3 (May 1994), 26–41.

[29] S. Notra, M. Siddiqi, H. Habibi Gharakheili, V. Sivaraman, and R. Boreli. 2014. An Experimental Study of Security and Privacy Risks with Emerging Household Appliances. In *IEEE Conference on Communications and Network Security*. 79–84.

[30] A. Panchenko, F. Lanze, A. Zinnen, M. Henze, J. Pennekamp, K. Wehrle, and T. Engel. 2016. Website Fingerprinting at Internet Scale. In *Proceedings of the 23rd Annual Network and Distributed System Security Symposium (NDSS)*.

[31] M. Patton, E. Gross, R. Chinn, S. Forbis, L. Walker, and H. Chen. 2014. Uninvited Connections: A Study of Vulnerable Devices on the Internet of Things (IoT). In *IEEE Joint Intelligence and Security Informatics Conference*. 232–235.

[32] F. Rath, J. Krude, J. Rüth, D. Schemmel, O. Hohlfeld, J. Á. Bitsch, and K. Wehrle. 2017. SymPerf: Predicting Network Function Performance. In *Proceedings of the ACM SIGCOMM Posters and Demos*.

[33] R. J. Robles and T. Kim. 2010. Applications, Systems and Methods in Smart Home Tech.: A Review. *International Journal of Advanced Science And Technology* 15 (Feb. 2010).

[34] S. Seeber and G. D. Rodosek. 2014. Improving Network Security Through SDN in Cloud Scenarios. In *10th International Conference on Network and Service Management (CNSM) and Workshop*. 376–381.

[35] A. Sivanathan, D. Sherratt, H. H. Gharakheili, V. Sivaraman, and A. Vishwanath. 2016. Low-Cost Flow-Based Security Solutions for Smart-Home IoT Devices. In *International Conference on Advanced Networks and Telecommunications Systems (ANTS)*. IEEE.

[36] V. Sivaraman, D. Chan, D. Earl, and R. Boreli. 2016. Smart-Phones Attacking Smart-Homes. In *Proceedings of the 9th Conference on Security & Privacy in Wireless and Mobile Networks (WiSec)*. ACM, 195–200.

[37] V. Sivaraman, H. H. Gharakheili, A. Vishwanath, R. Boreli, and O. Mehani. 2015. Network-Level Security and Privacy Control for Smart-Home IoT Devices. In *IEEE 11th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*. 163–167.

[38] The Open Networking Foundation. 2018. OpenFlow. https://www.opennetworking.org/technical-communities/areas/specification/open-datapath/. (2018).

[39] C. C. Tu, J. Stringer, and J. Pettit. 2017. Building an Extensible Open vSwitch Datapath. *SIGOPS Oper. Syst. Rev.* 51, 1 (Sept. 2017), 72–77.

[40] Q. Wang, W. U. Hassan, A. Bates, and C. Gunter. 2018. Fear and Logging in the Internet of Things. In *Proceedings of the Network and Distributed Systems Symposium (NDSS)*. Internet Society.

[41] J. Wurm, K. Hoang, O. Arias, A. R. Sadeghi, and Y. Jin. 2016. Security Analysis on Consumer and Industrial IoT Devices. In *Asia and South Pacific Design Automation Conference (ASP-DAC)*. 519–524.

[42] G. Xu, Y. Cao, Y. Ren, X. Li, and Z. Feng. 2017. Network Security Situation Awareness Based on Semantic Ontology and User-Defined Rules for Internet of Things. *IEEE Access* 5 (Aug. 2017), 21046–21056.

[43] I. Yaqoob, E. Ahmed, M. H. ur Rehman, A. I. A. Ahmed, M. A. Al-garadi, M. Imran, and M. Guizani. 2017. The Rise of Ransomware & Emerging Security Challenges in the IoT. *Computer Networks* 129 (Dec. 2017), 444–458.