Challenges in inferring privacy properties of smart devices: Towards scalable multi-vantage point testing methods

Aniketh Girish
IMDEA Networks Institute / Universidad Carlos III de Madrid

Vijay Prakash
New York University

Serge Egelman
ICSI / UC Berkeley

Joel Reardon
University of Calgary

Juan Tapiador
Universidad Carlos III de Madrid

Danny Yuxing Huang
New York University

Srdjan Matic
IMDEA Software Institute

Narseo Vallina-Rodriguez
IMDEA Networks Institute

ABSTRACT

The growth of the number of Internet of Things (IoT) devices in the home has introduced unprecedented challenges for preserving consumers' privacy. To enhance the transparency and trustworthiness of IoT products, industry actors, regulators and standardization bodies have proposed several certifications processes. While previous research developed methodologies to expose a wide range of harmful behaviors in smart home devices, the research community has not yet produced scalable methods to exhaustively detect information leakage in heterogeneous environments so that they can be effectively used for independent certification processes. Research questions such as *to what extent is it possible to automatically and independently validate high-level privacy properties* and *can we evaluate the correctness of privacy controls in a smart home by analyzing the home traffic?* remain open. This paper explores and discusses open research challenges in this complex domain and sketches a roadmap to build scalable methods for smart home privacy testing.

CCS CONCEPTS

• **Security and privacy → Network security**.

KEYWORDS

IoT, Privacy, Testing strategies, Measurement techniques

ACM Reference Format:
Aniketh Girish, Vijay Prakash, Serge Egelman, Joel Reardon, Juan Tapiador, Danny Yuxing Huang, Srdjan Matic, and Narseo Vallina-Rodriguez. 2022. Challenges in inferring privacy properties of smart devices: Towards scalable multi-vantage point testing methods. In *CoNEXT Student Workshop 2022 (CoNEXT-SW '22), December 9, 2022, Roma, Italy.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3565477.3569148

1 INTRODUCTION

The proliferation of consumer IoT devices in households has come at the cost of users' privacy [10]. Organizations like NIST [8] and ioXt [4] have proposed certification programs to independently validate the security and privacy attributes of an IoT product. However, there is a technical gap between the expectations of IoT certification frameworks and the actual capabilities of modern IoT testing tools for examining such properties.

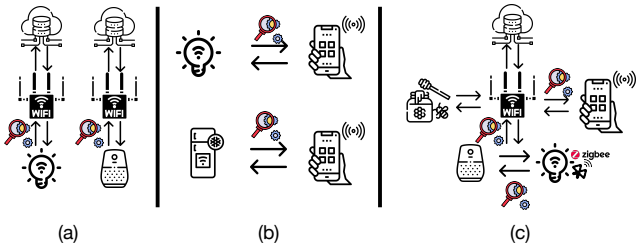

Figure 1: Different approaches to infer privacy properties in IoT devices ((a) in-path observer [10] and (b) ad-hoc methods [2]) vs. (c) our multi-vantage point solution.

Black-box testing has become the de-facto solution for independent IoT product certification due to the closed nature of most products. However, black-box testing involves inferring system properties from potentially incomplete signals which are only triggered when the system under analysis receives the right input [1]. Prior research has not yet produced methodologies to effectively and exhaustively test the privacy of any arbitrary smart home device when is co-located with other devices. As Figures 1.a and 1.b show, prior efforts tend to oversimplify the real nature of IoT devices and treat them as monolithic, deterministic and self-contained objects that just communicate with cloud services [2, 10], also using in-path vantage points with limited visibility over their payloads. But these assumptions are invalid as real smart home devices and platforms offer complex APIs to interact with other devices and software, and integrate rich technologies from a large supply chain of providers, including third-party AI libraries. As a result, many privacy-intrusive behaviors are not detected with existing methods.

Therefore, *are current black-box testing methods enough to certify the privacy and security properties of IoT products?* And *what are the theoretical, technical and methodological barriers that must be solved?* This paper aims to fill this knowledge gap. We discuss four research challenges that must be addressed to enable effective muti-vantage point methods for testing privacy properties of IoT products, as shown in Figure 1.c.

2 RESEARCH CHALLENGES

A review of works in the area of IoT privacy analysis using black-box testing methods reveals four open research problems. To motivate

CoNEXT-SW '22, December 9, 2022, Roma, Italy

ACM ISBN 978-1-4503-9937-1/22/12.
https://doi.org/10.1145/3565477.3569148

each challenge, we rely on prior work findings and empirical data from two complementary sources: (*i*) empirical behavior data gathered from our purpose-built lab with 16 IoT devices using a WiFi AP, UPnP, HTTP(s), mDNS and Telnet *honeypots*, and Frida-based scripts for dynamic testing of Smart TVs and companion smartphone apps; and (*ii*) crowd-sourced traffic data associated with over 10k IoT devices captured via ARP spoofing and active mDNS and UPnP/SSDP network scans provided by IoT Inspector [3].

Challenge 1: Handling diverse interfaces, protocols and dual purpose. Home IoT devices support a wide range of network interfaces and short-range protocols like WiFi, Bluetooth or ZigBee which must be actively monitored to detect (even unintended) privacy leaks. In fact, privacy-intrusive behaviors might occur in unexpected or difficult-to-monitor channels. For example, IoT devices locate peer devices through unique device IDs (e.g., MAC addresses) gathered through network scans. However, these IDs can be abused to fingerprint and geolocate households, assess their socio-economic status, and even infer social structures by monitoring roaming devices across networks. IoT Inspector data confirms the high prevalence of PII data in UPnP/SSDP flows within devices in the home network. Using this dataset, we find that over 38% of the labeled devices share unique device names or serial numbers, and over 7% share the MAC address. Our IoT and mobile app instrumentation also revealed that this data can be used by mobile apps for fingerprinting the household. Even third-party monetization libraries like the "Coelib" SDK which is present in mobile apps use multicast UPnP and Netbios scans for tracking purposes. Effectively tackling such risks requires the ability to correctly detect when a network interface or protocol of gets misused. Monitoring the right vantage points (*e.g.,* monitoring ARP cache of WiFi AP) to analyze the communication patterns across any set of IoT devices and applications can detect these unexpected privacy risks, but it imposes many challenges in terms of scalability and coverage.

Challenge 2: Automatic input generation. Due the presence of diverse input modalities and interfaces, automatically generating inputs for consumer IoT products is extremely challenging. For example, smart bulbs do not have a UI except for users' commands received from companion apps, automation platforms (e.g., IFTTT) or from IoT hubs. Wihout human interventions, NLP-based techniques are insufficient for exhaustively interacting with smart speakers [6]. Similarly, UI fuzzers and monkeys used to navigate through the UIs of smart TVs cannot trigger many software behaviors [7, 11]. We build on prior work by prototyping a new type of *smart "monkey"* for IoT products that simulates real user interactions based on the state transition model of an app. This is component is then integrated by a smart network honeypot that passively captures traffic from device scans and requests. While our preliminary results are promissing, it is crucial to gain a good understanding of the user and network interfaces of both IoT devices and apps in order to increase its ability to capture additional behaviors. Yet, our honeypot allowed us to discover the use of 1-byte long UDP packets on random ports in order to receive ICMP packets from co-located devices (and apps) with open ports; Amazon devices (e.g., Echo dot and Fire TV) communicate over UDP:55444, possibly for spatial perception or synchronizing time [12]. We plan to extend our ongoing methods by (*i*) integrating automatic voice command generation for voice assistants; and (*ii*) leveraging reinforcement learning methods to develop a network-level fuzzer that can virtually emulate any arbitrary smart devices or services without relying on human inputs and without deploying their physical counterpart.

Challenge 3: Increasing adoption of anti-testing methods. The increased in use of security measures such as encrypting traffic using TLS [9] or anti-testing methods provides significant benefits to consumers' security. However, these techniques have direct implications on the ability to monitor behaviors of smart devices. Except Smart TV platforms, most other IoT platforms and products are closed-source and they do not provide practical debugging capabilities to inspect behaviors. Researchers have explored solutions that instrument the Alexa and emulate it on Raspberry Pi which (partially) mimics the features and capabilities of a real device [5]. However, without built-in debugging capabilities, creating valid countermeasures for anti-testing of every device by just relying on low-level source-code instrumentation is a huge engineering and research effort that may be difficult to maintain and evolve as new releases hit the market.

Challenge 4: Simulating and measuring user consent. Modern software must comply with new regulatory frameworks like GDPR and CCPA. For instance, both require developers to obtain informed consent from users and state the purpose of data collection before collecting any personal data. However, as described in *Challenge 2*, there are no reliable methods for automatically interacting and validating whether consent was obtained when personal data was collected on the device. Prior work approached this problem in mobile and web technologies by detecting inconsistencies between developer claims in privacy policies (extracted with NLP methods) and then finding contradictions in software runtime behavior [6]. However, contextualizing software behaviors from a privacy perspective in the presence of IoT devices setups that support new interfaces (e.g., VR, voice commands or BCI), and the dissemination of personal data across devices and apps is more challenging. Solving this problem requires multidisciplinary expertise in areas such as NLP, AI and software testing.

3 DISCUSSION AND FUTURE WORK

Modern IoT devices are not isolated entities. This makes traditional black-box testing approaches fundamentally inadequate to capture the dynamic, unexpected and privacy-intrusive behaviors of interconnected IoT devices. This paper describes and discusses open research challenges that prevents the research community from developing robust and scalable methods for independent and exhaustive testing of IoT devices.

ACKNOWLEDGMENTS

This project has been funded by the EU's H2020 Program (TRUST aWARE Project, Grant Agreement No. 101021377), the Spanish Ministry of Science (ODIO Project, PID2019-111429RB-C22), the Atracción de Talento grant (Ref. 2020-T2/TIC-20184) funded by Madrid regional government, the SCUM Project (RTI2018-102043-B-I00) MCIN/AEI/10.13039/ 501100011033/ERDF and comunidad de Madrid (P2018/TCS-4566). The opinions, findings, and conclusions or recommendations expressed are those of the authors and do not necessarily reflect those of any of the funders.

REFERENCES

[1] Mohammad Torabi Dashti and David Basin. 2020. A Theory of Black-Box Tests. *CoRR* abs/2006.10387 (2020). http://arxiv.org/abs/1708.05044

[2] Dennis Giese and Guevara Noubir. 2021. Amazon Echo Dot or the Reverberating Secrets of IoT Devices. In *Proceedings of the 14th ACM Conference on Security and Privacy in Wireless and Mobile Networks*.

[3] Danny Yuxing Huang, Noah J. Apthorpe, Frank Li, Gunes Acar, and Nick Feamster. 2020. IoT Inspector: Crowdsourcing Labeled Network Traffic from Smart Home Devices at Scale. *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT)* (2020).

[4] ioXt. 2020. The ioXt Security Pledge. https://www.ioxtalliance.org/the-pledge.

[5] Umar Iqbal, Pouneh Nikkhah Bahrami, Rahmadi Trimananda, Hao Cui, Alexander Gamero-Garrido, Daniel Dubois, David Choffnes, Athina Markopoulou, Franziska Roesner, and Zubair Shafiq. 2022. Your Echos are Heard: Tracking, Profiling, and Ad Targeting in the Amazon Smart Speaker Ecosystem. http://arxiv.org/abs/2204.10920

[6] Tu Le, Danny Yuxing Huang, Noah Apthorpe, and Yuan Tian. 2022. SkillBot: Identifying Risky Content for Children in Alexa Skills. In *ACM Transactions on Internet Technology*.

[7] Hooman Mohajeri Moghaddam, Gunes Acar, Ben Burgess, Arunesh Mathur, Danny Yuxing Huang, Nick Feamster, Edward W. Felten, Prateek Mittal, and Arvind Narayanan. 2019. Watching You Watch: The Tracking Ecosystem of Over-the-Top TV Streaming Devices. In *Conference on Computer and Communications Security (CCS)*.

[8] National Institute of Standards and Technology (NIST). 2021. DRAFT Baseline Security Criteria for Consumer IoT Devices. https://www.nist.gov/system/files/documents/2021/08/31/IoT%20White%20Paper%20-%20Final%202021-08-31.pdf.

[9] Muhammad Talha Paracha, Daniel J. Dubois, Narseo Vallina-Rodriguez, and David R. Choffnes. 2021. IoTLS: Understanding TLS Usage in Consumer IoT Devices. In *Proceedings of the Internet Measurement Conference (IMC)*.

[10] Jingjing Ren, Daniel J. Dubois, David R. Choffnes, Anna Maria Mandalari, Roman Kolcun, and Hamed Haddadi. 2019. Information Exposure From Consumer IoT Devices: A Multidimensional, Network-Informed Measurement Approach. In *Proceedings of the Internet Measurement Conference (IMC)*.

[11] Irwin Reyes, Primal Wijesekera, Joel Reardon, Amit Elazari Bar On, Abbas Razaghpanah, Narseo Vallina-Rodriguez, and Serge Egelman. 2018. "Won't Somebody Think of the Children?" Examining COPPA Compliance at Scale. *Proceedings on Privacy Enhancing Technologies (PoPETs)*.

[12] Tom Warren. 2018. Amazon is making it easier for all Alexa devices to work together. https://www.theverge.com/2018/7/25/17613832/amazon-alexa-echo-spatial-perception.