Towards automated privacy risk assessments in IoT systems

Milan Markovic
University of Aberdeen
Aberdeen, UK
milan.markovic@abdn.ac.uk

Waqar Asif
City, University of London
London, UK
waqar.asif@city.ac.uk

David Corsar
University of Aberdeen
Aberdeen, UK

Naomi Jacobs
University of Aberdeen
Aberdeen, UK
naomi.jacobs@abdn.ac.uk

Peter Edwards
University of Aberdeen
Aberdeen, UK
p.edwards@abdn.ac.uk

Muttukrishnan Rajarajan
City, University of London
London, UK
r.muttukrishnan@city.ac.uk

Caitlin Cottrill
University of Aberdeen
Aberdeen, UK
c.cottrill@abdn.ac.uk

ABSTRACT

Internet of Things (IoT) systems can often pose risk to users' privacy via disclosure of personal information to third parties. In this paper, we argue that to assess privacy risks associated with IoT systems, an automated solution is required due to the increasing pervasiveness and complexity of deployed IoT systems. We propose requirements for an automated privacy risk assessment service and outline our future plans for realising such a solution.

CCS CONCEPTS

• **Security and privacy** → **Domain-specific security and privacy architectures**; • **Computing methodologies** → *Semantic networks*; • **Computer systems organization** → *Sensor networks*;

KEYWORDS

Internet of Things, Privacy Risk , Provenance, Linked Data.

ACM Reference Format:
Milan Markovic, Waqar Asif, David Corsar, Naomi Jacobs, Peter Edwards, Muttukrishnan Rajarajan, and Caitlin Cottrill. 2018. Towards automated privacy risk assessments in IoT systems . In *M4IOT '18: Workshop on Middleware and Applications for the Internet of Things, December 10–11, 2018, Rennes, France.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

The Internet of Things (IoT) refers to a collection of heterogeneous devices and applications networked using the Internet and deployed to observe and actuate various aspects of physical environments. IoT systems can consist of sensors, actuators, gateways, cloud-based message brokers, and various applications supporting data management, analysis, and visualisation. Such systems can quickly become large distributed ecosystems with complex data sharing patterns. As a result it might be difficult for an individual to maintain their privacy by enacting "the power to selectively reveal oneself to the world" [4] when they may be unaware of the data being collected about them in an IoT context and how it is being used.

The General Data Protection Regulation (GDPR) dictates that the use and collection of data about an individual for purposes other than provision of direct services (e.g. research & development, business modelling, etc.) is allowable only for non-personal information, in order to maintain user privacy. As a result, any data collected by an IoT system has to be evaluated for the presence of identifiers[1] and quasi-identifiers[2]. When such identifiers are collected, anonymisation processes should be used, in which all identifiers need to be removed and all quasi-identifiers need to go through rigorous procedures to ensure user privacy [1].

In the IoT context, the risk of information privacy violation can be high due to the way IoT systems operate. Personal data can be "leaked", for example, via device tampering, intercepting of messages transferred on unsecured networks, and gaining access to the stored data [1]. Data can also be shared with third parties without the knowledge of the user or without their consent. We argue that due to the complexity and widespread use of IoT in everyday environments, it is unfeasible to expect an individual to assess potential privacy risks for every IoT system they encounter. To address this, we propose a privacy risk assessment middleware service to support user applications by informing them about privacy risks associated with IoT systems.

In this paper, we outline the service input requirements (i.e. information about the IoT systems, user risk preferences, etc.) and propose an approach for realising such a service. The remainder of the paper is structured as follows: Section 2 introduces a use case scenario to highlight privacy risks faced by IoT systems; Section 3 discusses different types of privacy risks and knowledge requirements to support assessment of these; Section 4 outlines a proposed architecture of an automated privacy assessment service; and Section 5 concludes with a discussion of our future plans to realise such a service.

[1] Information that can be directly related to an individual, such as name, email address and contact number.
[2] Information parameters that have no meaning on their own but if observed in correlation with other quasi identifiers can reveal the actual identity of an individual; for example age, location and ethnicity.

M4IOT '18, December 10–11, 2018, Rennes, France
2018. ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

2 USE CASE: IOT KETTLE

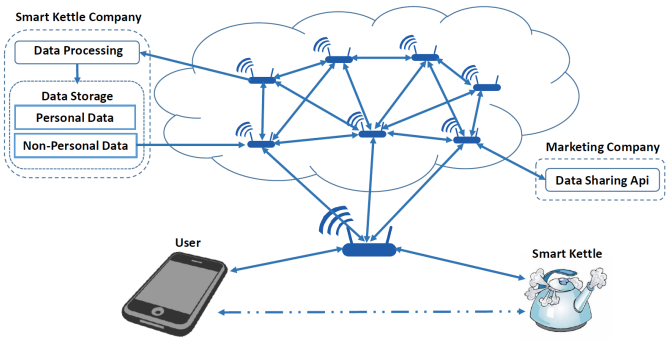

Figure 1: An illustration of main components of the smart kettle IoT infrastructure.

To illustrate risks associated with sharing of personal information in IoT systems, let us consider the case of smart kettles as an example. A common household appliance (the kettle) has been extended with "smart" capabilities. Along with the "original" ability (boiling water), "smart" kettles are connected to the Internet allowing the operation of the kettle to be controlled via a smart phone app, the ability to set the temperature water is heated to based on a specific requirement (e.g. to prepare baby formula), and real-time monitoring of the water temperature. In order to provide these features, smart kettles are typically supported by a cloud-based IoT architecture [5] which handles the exchange of data between the kettle, cloud services operated by the manufacturer (or a company acting on their behalf), and the app (users). Figure 1 illustrates such an architecture. The *user* interacts with the kettle via a smart phone app which is backed by an API provided by the *Smart Kettle Company* - i.e. a company responsible for delivery of smart kettle features. The *Smart Kettle Company* runs a cloud service that enables collection of data from the kettle (e.g. current water temperature) and control of the device (e.g. starting the boiling cycle). It also provides the user's mobile app which requires users to provide personal information when registering the kettle with the app. Both this personal data (name, email, registered kettle id) and further data (location, water temperatures, time of starting/stopping the boiling cycle) are then stored on company's servers. In such a scenario, the more entities involved in the information flow and/or storage, the higher the risk to a kettle owner's privacy. An adversary can, for example, violate their privacy either by tampering with the kettle physically or by launching an eavesdropping attack to gain personal information such as device id or location and use this to profile an individual. Data can also be shared without the owner's knowledge with, for example, a *Marketing Company* for targeted marketing. A simple device such as a kettle can then reveal vital information such as the work pattern of individuals in the household, or can highlight the presence of an infant based on the "baby bottle" feature usage pattern.

3 AUTOMATED ASSESSMENT OF PRIVACY RISKS

The smart kettle scenario introduced in the previous section is a typical example of W3C Web of Things IoT architecture for cloud ready devices [5]. It highlights a number of potential privacy risks due to the transfer and storage of personal data (name, email, associated kettles) and quasi-personal data (kettle id with event timestamps). We will now discuss the different types of privacy risks and the information about IoT systems that is required in order to assess them.

3.1 Privacy Risks in IoT Systems

The privacy risks associated with the use of an IoT device can emerge from three parts of the device's setup: the device itself, the communication channel, and the data storage location [1].

Most IoT devices such as the smart kettle are small in size and are physically accessible to the device user. These devices are designed and developed by various vendors thus posing a large variety of privacy risks. The small size and accessible nature increases the possibility of device tampering by ill intent of the user (e.g. to tamper with the device data) or theft by an external adversary. The large number of vendors combined with the ease of purchasing a device increases the chances of coming across a device that does not comply with the latest privacy policies such as the GDPR or relies on outdated privacy definitions resulting in privacy violations even through legally shared data. In such context, a privacy definition refers to the combination of encryption and anonymisation approaches that are used to ensure user privacy. The lack of an updated privacy definition results in a high probability of a privacy violation through the communication channel. These IoT devices relay information to the backhaul network via the Internet thus relying on multiple intermediate nodes. The large number of intermediate nodes makes it possible for an adversary to overhear the communication via eavesdropping thus obtaining personal information without a user's consent. IoT device manufacturers manage their devices with the help of data processing and storage facilities. They gather information from all devices, process it to ensure privacy and then make that information available to the device user. Manufacturers can either host information on their own servers or can hire third party companies for such a purpose. In either case, it is crucial to monitor who has access to the stored dataset. Conventionally, an anonymised version of this stored dataset is shared with third party companies for the purpose of analysis and research. The amount of risk that this shared data poses depends heavily on the privacy definition that is in place. An anonymised version of a dataset can be correlated to data obtained from secondary sources if the data is not carefully anonymised. Similarly, a third party company hosting such a dataset can exploit user's privacy by illegally sharing this data with data analysts thus posing a high privacy risk [1]. A detailed division and explanation of these potential risks can be found in [1].

A user can face a risk of privacy violation due to any of the aforementioned reasons. Therefore, in order to evaluate the potential privacy risk, one needs to take into account all possible venues that an adversary can exploit. Building on our previous work [1], we

propose to exploit this approach also in the design of the automated risk assessment service which is discussed further in Section 4.

3.2 Enabling Transparency of IoT Systems to Support Risk Assessment

The risks discussed above can originate from several locations, such as the kettle, the app, transfer of data between the app and kettle, kettle and cloud services, and sharing of data stored by the company. For computational processes to be capable of automatically identifying such privacy risks, it is necessary to develop robust auditing mechanisms, which pose a challenge in a distributed IoT environment [9]. We propose to address this challenge by recording structured, machine readable descriptions of the various IoT concepts involved in the deployment/architecture. These descriptions would include information describing IoT devices, how data is being acquired and used, agents involved in data processing, and so on. Semantic technologies such as ontologies and linked data[3] provide a method to define machine processable model of domain concepts, such as IoT devices, and their inter-relationships, which can be shared and reused across different IoT deployments and reasoning applications, such as the privacy assessment service proposed by this paper.

Ontologies are expected to play an important part of future IoT systems [10]. One such ontology is the Semantic Sensor Network Ontology (SSNO), a W3C recommendation [6] which provides formalisms to describe sensors and related concepts in domains such as the Internet of Things. SSNO can be used to describe systems, such as the smart kettle, in terms of the sensor, actuators, samplers, and other types of system that it is composed of. The capabilities, survival, and operating ranges of each system can also be specified. For example, the smart kettle system includes a temperature sensor which has a survival and operating range of at least between 0 and 100 °C, can measure temperature with an accuracy of +/- 1 °C with a frequency of at least 1 second between observations. SSNO can also describe the observations made by the temperature sensor, in terms of their type (observations of temperature), the feature of interest (i.e. the thing being observed - the water in the kettle), the property of that that is being observed (i.e. temperature), and the observation result value (e.g. 99), unit of measurement (°C), and time that the observation was made.

SSNO can also describe the deployment of the smart kettle, in terms of where it has been installed. Further details of the deployment can be described using the SSN System Deployment Provenance Ontology (SDPO) [2]. SDPO extends the PROV W3C recommendation [8] for documenting provenance to support describing any of the activities (e.g. installation, system setup, device purchase, registration, and maintenance operations) that have been performed before (or during) the deployment, the agents (people, software agents, or organisations) associated with those activities (and their role - for example, a company selling the device, or the homeowner/user setting it up) and other entities (things) used (e.g. the app used to register the kettle, the worktop on which it is installed). SDPO also allows the expected behaviour of the device to be described in terms of the process that it will enact during a deployment using extensions of the P-Plan ontology [3]. For example, the kettle has a plan to heat the water in response to a user's instruction, until the temperature sensor observes a certain temperature has been reached.

P-PLAN and PROV have been also extended in another of our recently developed ontologies MQTT-PLAN [7] designed to enable recording provenance generated by MQTT[4] message brokers. The ontology can be used to document the intended and actual behaviour of message brokers. This could then support audit of message distribution (e.g. to which agents was a message forwarded). In addition, such provenance could be used to discover issues relating to malfunctions, misconfigurations or the limited capabilities of message brokers (e.g. not detecting abnormal behaviour such as repeated failed authentication attempts). For example, such records can reveal a malicious agent repetitively attempting to request forwarding of messages under topics (e.g. the kettle usage) without appropriate authorisation.

3.3 Knowledge Gaps for Privacy Risk Assessment

As we have outlined in the previous section, several aspects of IoT systems can be described using machine processable semantic annotations. However, in order for a middleware service to automatically reason about and identify potential privacy risks, we argue that the following must also be modelled:

- Conceptualisation of the planned agency within IoT systems - i.e. representations of agents in relation to their intended responsibilities within a plan.
- Concepts enabling classification and tracking of various types of personal and quasi-personal data within an IoT system (both planned and actual).
- Concepts enabling classification and tracking of anonymisation techniques applied to data at any point in the IoT system (both planned and actual).
- Concepts describing cloud-based components that deal with processes such as storing and analysing of data.
- Concepts for linking descriptions of intended policies (e.g., privacy, security, user policies) and constraints applied to data and processes in IoT systems and provenance of compliance with these.
- Concepts representing results of the privacy risk assessment and their relationships to other concepts used to evaluate such risk.

4 PRIVACY RISK ASSESSMENT SERVICE

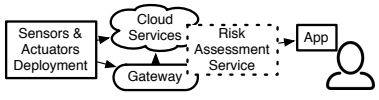

Figure 2: An illustration of the position of the privacy risk assessment service in the IoT architecture.

[3]https://www.w3.org/standards/semanticweb/data

[4]MQTT is a publish/subscribe messaging transport protocol for a client-server communication. The protocol specifies a set of control packets that govern the communication between the client and the message broker residing on a server.

The computational power required to perform privacy risk assessments typically cannot be delivered by edge devices such as sensors and actuators. Therefore, such service would exist as part of a gateway or as a cloud-based component in an IoT system (Figure 2). Figure 3 illustrates a proposed architecture of such a service.

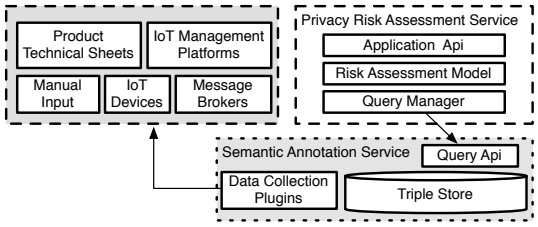

Figure 3: An illustration of a high level architecture of the proposed risk assessment service.

The semantic annotations stored in the triple store include descriptions of devices, data flow structures, types of data exchange, and provenance describing the planned and the actual behaviour of an IoT system are obtained from a range of heterogeneous sources. These might include IoT *management platforms* (e.g. NodeRed[5], Kura[6], OM2M[7]), *message brokers* (e.g. Mosquitto[8]), individual *IoT devices* (e.g. sensors), hardware specification sheets / data sheets (e.g. Texas Instruments website[9]), and manual inputs (e.g. from installation engineers, users, etc.). The semantic annotations are modelled using extensions of SSNO to describe the IoT devices and domain specific extensions of PROV and P-PLAN (e.g. SDPO and MQTT-PLAN) to document systems' provenance as discussed in section 3.2. Annotations are stored in a semantic data store (*triple store*) and can be queried via an API endpoint using SPARQL query syntax[10]. Results of such queries represent various information about an IoT system and serve as input parameters of a privacy risk assessment service.

The privacy assessment service evaluates the input information based on a risk model which considers three parts of the network that can pose a risk: the device, the communication channel and the data storage location. To realise this component, we propose to implement an attack tree model introduced in our previous work [1] with any potential modifications to accommodate new input parameters resulting from rich semantic descriptions of IoT systems. The attack tree model makes use of the logical AND/OR gate structure to form a hierarchical structure, where the parent node, referred to as the goal of the tree, is the risk of privacy violation, the intermediate nodes are referred to as sub-goals and the leaf nodes are individual privacy risk attacks such as risk of device tampering. Each independent path in the risk tree becomes a separate scenario and each of these scenarios has a different probability of occurrence. The probability of occurrence of a scenario depends on factors such as the attack impact, the cost of an attack, the technical difficulty and the probability of being discovered. For instance, the probability of having a device tampering attack is greater for a smart kettle than a smart meter as the cost of an attack and the technical difficulty of tampering with a smart meter is greater than that of a kettle. This means that in order to assess the privacy risk for an IoT device, one needs to consider all device parameters, the communication framework and the data storage protocols. This would help analyse the probability of occurrence of each sub-goal and the leaf node in an attack tree model thus resulting in better approximation of the risk associated to that device. The results of the assessment can be made available via application API to enable development of interfaces informing users about privacy risks. Such results can then be presented in combination with the semantic annotations used by the service, for example, to inform about the coverage and type of information about an IoT system the service had access to when assessing risk.

[5]https://nodered.org
[6]https://www.eclipse.org/kura/
[7]http://www.eclipse.org/om2m/
[8]https://mosquitto.org
[9]http://www.ti.com/
[10]https://www.w3.org/TR/sparql11-overview/

5 FUTURE WORK

As part of our future work we aim to extend the ontologies discussed in Section 3.2 according to the requirements outlined in Section 3.3. In addition, we plan to develop a prototype realisation of the discussed service and perform evaluation using real world IoT deployments. For this we will ustilise our community testbed environment established as part of the TrustLens[11] project that is investigating issues related to transparency and accountability of IoT systems as well as empowerment of the users via providing control over these systems.

ACKNOWLEDGMENTS

The work described here was funded by the award made by the RCUK Digital Economy programme to the University of Aberdeen (EP/N028074/1) and City, University of London (EP/N028155/1).

REFERENCES

[1] W. Asif, I. G. Ray, and M. Rajarajan. 2018. An attack Tree Based Risk Evaluation Approach for The Internet of Things. In *Proceedings of the 8th International Conference on the Internet of Things*. ACM, Santa Barbara, USA.
[2] D. Corsar, M. Markovic, and P. Edwards. 2018. Capturing the Provenance of Internet of Things Deployments. In *Proceedings of the 7th International Provenance & Annotation Workshop-IPAW*. Springer, London, UK, 196–199.
[3] D. Garijo and Y. Gil. 2012. Augmenting PROV with Plans in P-PLAN: Scientific Processes as Linked Data. In *Proceedings of the Second International Workshop on Linked Science 2012 - Tackling Big Data*. CEUR-WS, Boston,USA.
[4] E. Hughes. 1993. A cypherpunk's manifesto. (1993). http://www. activism. net/cypherpunk/manifesto. html.
[5] K. Kazuo, K. Matthias, and D. Uday. 2017. *Web of Things (WoT) Architecture*. W3C Working Draft. W3C. https://www.w3.org/TR/2017/WD-wot-architecture-20170914/.
[6] M. Lefrançois, S. Cox, K. Taylor, A. Haller, K. Janowicz, and D. Le Phuoc. 2017. *Semantic Sensor Network Ontology*. W3C Recommendation. W3C. https://www.w3.org/TR/2017/REC-vocab-ssn-20171019/.
[7] M. Markovic, D. Corsar, and P. Edwards. 2018. Towards Transparency of IoT Message Brokers. In *Proceedings of the 7th International Provenance & Annotation Workshop-IPAW*. Springer, London, UK, 200–203.
[8] L. Moreau, P. Groth, J. Cheney, T. Lebo, and S. Miles. 2015. The rationale of PROV. *Web Semantics: Science, Services and Agents on the World Wide Web* 35 (2015), 235–257.
[9] J. Singh, T. Pasquier, J. Bacon, J. Powles, R. Diaconu, and D. Eyers. 2016. Big ideas paper: Policy-driven middleware for a legally-compliant Internet of Things. In *Proceedings of the 17th International Middleware Conference*. ACM, Trento, Italy.
[10] I. Toma, E. Simperl, and G. Hench. 2009. A joint roadmap for semantic technologies and the internet of things. In *Proceedings of the Third STI Roadmapping Workshop*. Crete, Greece, 140–53.

[11]http://trustlens.org