Interrogating Capabilities of IoT Devices

Stanislav Beran, Edoardo Pignotti(✉), and Peter Edwards

Computing Science and dot.rural Digital Economy Hub, University of Aberdeen, Aberdeen AB24 5UA, UK
{s.beran,e.pignotti,p.edwards}@abdn.ac.uk

Abstract. In this demo we present the *Trusted Tiny Things* system that can be used to interrogate Internet of Things (IoT) devices and present users with information about their characteristics and capabilities. The system consists of a mobile application used to retrieve information about IoT devices supported by RESTful web services. In order to infer IoT device capabilities our services perform reasoning over the provenance of devices characterised using an extension of the PROV-O ontology. In this demo we illustrate the use of the system with two distinct IoT devices: an NFC tag used at bus stops to provide a means to access real-time bus timetables, and a blackbox device installed into vehicles by insurance companies to track driving behaviour.

Keywords: Internet of things · Provenance · Transparency

1 Introduction

The *Trusted Tiny Things project*[1] is exploring how semantic technologies can make Internet of Things (IoT) devices more transparent to users. IoT devices now routinely gather, analyse and manipulate data from their surroundings; they are also capable of exchanging such data with other devices and services by means of M2M (Machine to Machine) communications. The need for transparency in the IoT domain is seen as crucial in order to ensure the legitimacy of activities performed by devices, but also to increase security and privacy [1]. Certain operations associated with IoT devices may be deemed undesirable by users (e.g. third-party data sharing, consumption of personal data), and therefore users should be made aware of such capabilities. In this paper we argue that by publishing information about IoT devices such as manufacturer, owner, device type) according to the linked data principles [2] and by capturing their provenance (e.g. services, owners, organisations, etc.), it is possible to make capabilities of IoT devices more transparent.

We are investigating these issues via two user scenarios. The first of these explores the use of NFC tags attached to timetables at bus stops in Aberdeenshire, UK. A user with an NFC enabled phone can scan such tags to access a

[1] This research is supported by the UK Research Councils' Digital Economy IT as a Utility Network+ (EP/K003569/1) and the dot.rural Digital Economy Hub (EP/G066051/1).

B. Ludäscher and B. Plale (Eds.): IPAW 2014, LNCS 8628, pp. 197–202, 2015.
DOI: 10.1007/978-3-319-16462-5_15

real-time bus timetable via the phone's web browser. Users may expect that the service is operated by Aberdeenshire Council, but in fact it is run by an external IT solutions provider. As part of offering the service, this third party organisation collects data from the smartphone (e.g. IP address, type of smartphone device). Our second user scenario investigates the use of in-car blackboxes, which are being installed into vehicles by insurance companies. These devices are used to track driver's behaviour in order to tailor insurance premiums to individuals. The devices continuously collect data (e.g. GPS location, acceleration, driving patterns, etc.) and connect to a third party service that collects the data on behalf of the insurance company. In this scenario the service could change over time. For example, a new organisation (e.g. a car manufacturer analysing engine management data) could be allowed to use the data generated by the sensors.

2 Semantic Framework

In order to inform the design of a semantic framework for IoT devices we have conducted three participatory design events involving a total of 14 participants with different technological backgrounds. Participants were asked to discuss issues surrounding the capabilities of IoT devices. Questions were posed such as: *What do you think are the capabilities of this device?* and *What kind of capabilities would you want to be aware of before interacting with this kind of device?*. We have developed an OWL ontology[2] (illustrated in Fig. 1) to link physical entities (*iota:PhysicalEntity*) with their IoT components (*iota:Device*[3]) using concepts derived from a model created as part of the Internet of Things Architecture (IoTa) project[4]. This allows us to identify those IoT devices and their virtual representations (*iota:VirtualEntity*) so we can analyse their characteristics and capabilities. The PROV-O [3] ontology is used as an upper ontology and allows us to characterise entities (data), activities (device processes and operations) and agents (either software or physical) associated with IoT devices and supporting services. For example, we can associate a particular device activity (e.g. location sensing) to the agent that initiated the operation (e.g. insurance company). Using PROV-O allows queries to be formulated such as: Who initiated the action? What entities have been used? When was a particular action executed? However, PROV-O on its own cannot answer questions such as: Why and for what purpose were the data used? Is the data confidential?

Guided by user requirements we have designed an ontology to support inferences about device capabilities using provenance described according to the PROV-O and IoTa ontology. We created a lightweight ontological model called T3[5] that provides annotations over provenance records. Using this model, we

[2] http://t3.abdn.ac.uk/ontologies/iota.owl.

[3] An artefact that provides an interface between the digital world and the physical world.

[4] http://www.iot-a.eu.

[5] http://t3.abdn.ac.uk/ontologies/t3.owl.

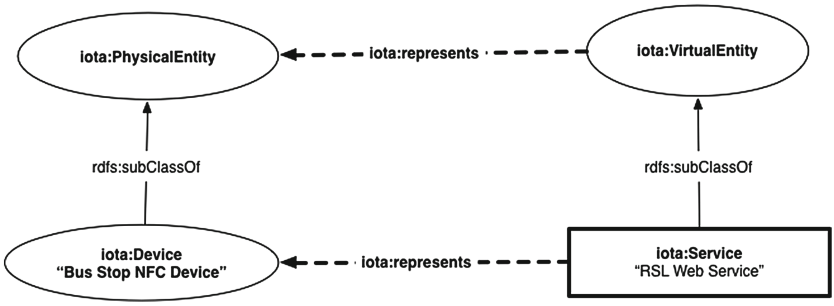

Fig. 1. An extract of the iota ontology representing relationships between a virtual entities and a physical entities in the internet of things

are able to annotate the qualified usage class (*prov:Usage*) with *ttt:purpose* to describe why a particular entity (data) is used by a specific activity.

When managing provenance of IoT devices it is not always possible to instrument devices and services to generate information about their usage and operation (retrospective provenance). In some cases, manufacturers can provide information on how devices are intended to operate (prospective provenance). In our framework we therefore make provision for both kinds of provenance. Our framework is also capable of distinguishing between direct capabilities (activities performed onboard the device) and indirect capabilities (activities performed by associated devices or services).

In order to infer the capabilities of IoT devices using our ontological framework we can associate rules to specific classes of *ttt:Capability*. We make use of the SPIN ontology[6] to support the use of SPARQL to specify rules and logical constraints necessary to reason about capabilities. The SPIN ontology allows SPARQL queries to be represented in RDF and associated to classes in an ontology using a pre-defined *spin:rule* property that can be used to specify inference rules using SPARQL CONSTRUCT, DELETE and INSERT statements. Figure 2 (top box) shows an example of such a rule for the *ttt:DataConsumption* class. The rule is designed to traverse a PROV-O provenance graph starting from an instance of an *iota:Device* and to identify activities that have used or generated entities classified as personal data. Once such activities have been identified the rule specifies how an annotation about the data consumption capability is generated, including a link to the agent responsible for the activity and the specific purpose. In this ontology we have also specified two rules that are used to determine what provenance has been used to infer a specific device capability. These rules make use of the *ttt:Follows* qualified relationship to distinguish between prospective and retrospective provenance and are illustrated in Fig. 2 (bottom left and bottom right boxes).

Participants during our design exercises highlighted the need to provide contact information about agents (individuals or organisations) responsible for

[6] http://spinrdf.org/spin.html.

```
CONSTRUCT {
   _:b0 a :DataConsumption .
   _:b0 :consumes ?data .
   _:b0 :consumer ?agent .
   _:b0 :purpose ?purposeDescription .
   ?this :isCapableOf _:b0 .
}
WHERE {
   ?virtualentity iota:represents ?this .
   ?activity (prov:qualifiedUsage)+ ?usage .
   ?usage prov:entity ?data .
   ?data a :PersonalData .
   ?usage :purpose ?purposeDescription .
   ?activity (prov:wasAssociatedWith)+/prov:actedOnBehalfOf ?agent .
   ?agent a foaf:Organization .
   NOT EXISTS {
      ?this :isCapableOf ?capability .
      ?capability :purpose ?purposeDescription .
      ?capability :consumer ?agent .
      ?capability :consumes ?data .
   } .
}
```

```
CONSTRUCT {
   ?device :prospectiveCapability ?capability .
}
WHERE {
   ?device prov:wasAttributedTo ?agent .
   ?agent :qualifiedFollow ?follow .
   ?follow :shouldGenerate ?bundle .
   ?device :isCapableOf ?capability .
   ?capability :consumes ?data .
   ?bundle :contains ?data .
   NOT EXISTS {
      ?device :prospectiveCapability ?capability .
   } .
}
```

```
CONSTRUCT {
   ?device :retrospectiverCapability ?capability .
}
WHERE {
   ?device prov:wasAttributedTo ?agent .
   ?bundle a prov:Bundle .
   ?bundle prov:wasAttributedTo ?agent .
   ?device :isCapableOf ?capability .
   ?capability :consumes ?data .
   ?bundle :contains ?data .
   NOT EXISTS {
      ?device :retrospectiveCapability ?capability .
   } .
}
```

Fig. 2. Example of device capability inference rule (top box) and two rules used to distinguish between prospective and retrospective provenance (bottom left and bottom right boxes).

certain devices and therefore we use the FOAF[7] ontology. The class *foaf: Organization* is defined as a subclass of *prov:Agent*. Figure 3 presents a visualisation of the device capabilities in a mobile app and the respective sample provenance graph taken from the bus stop scenario.

3 The Trusted Tiny Things System

In order to support our semantic framework we have developed a software infrastructure (see Fig. 4) that can be used to query, update and register IoT devices and to notify the user of any changes in the capabilities of a particular device. We store device data in an OpenRDF Sesame[8] triplestore. Additionally, we utilize a MySQL database server to store smartphone IDs (used to identify users) and accepted device capabilities. Our framework is composed of five core services, which are responsible for registering devices to our system, updating and synchronizing the provenance record, providing access to information, reasoning over the provenance record to infer capabilities, and notifying users about changes in device provenance. In order for a user to interact with the system, we have implemented an Android mobile application (Fig. 3), that is able to query

[7] http://www.foaf-project.org/.
[8] http://www.openrdf.org.

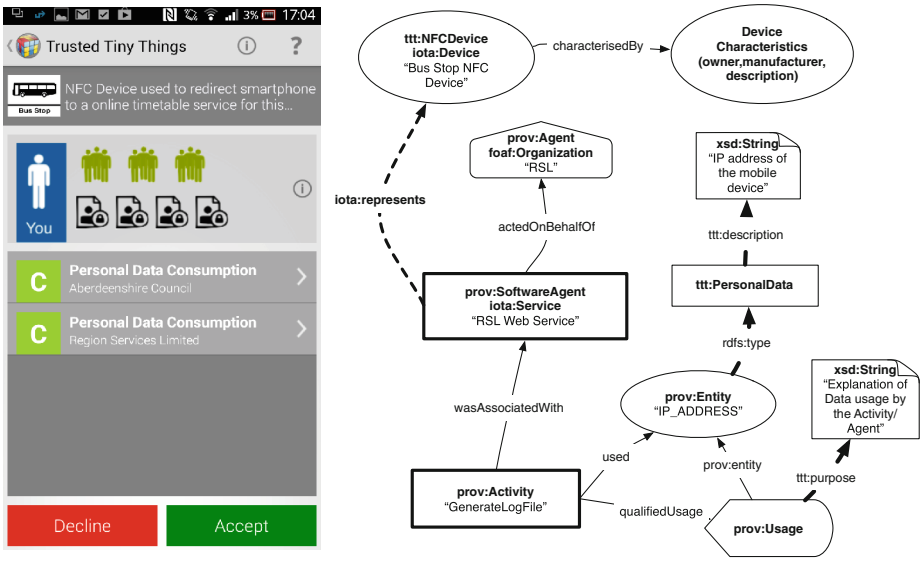

Fig. 3. Smartphone app showing the capabilities of a bus stop NFC tag (left) and an extract of the supporting provenance (right).

and visualise capabilities of IoT devices registered in our system and to notify users of changes in the provenance record. The application can be downloaded from the Google Play Store[9].

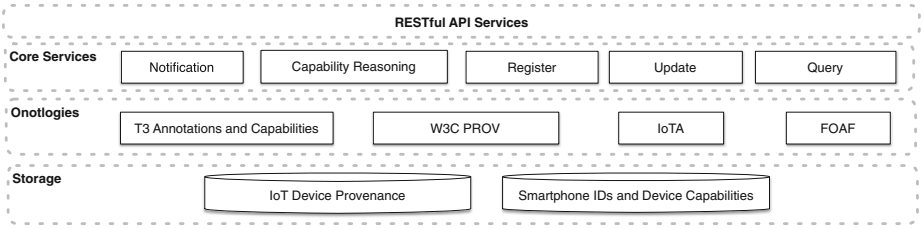

Fig. 4. Trusted Tiny Things System Architecture

The provenance-based approach for determining the capabilities of a device has certain advantages over similar compliance-based alternatives. In Google Play, for example, users are presented with a list of access permissions based on the capabilities of the app being installed. These permissions are determined only by the functionalities implemented in the app (e.g. use of the GPS sensor) disregarding how and why information is used and by whom. However, by using

[9] https://play.google.com/store/apps/details?id=uk.ac.abdn.t3.trustedtinythings.

a provenance-based approach, it is possible to define capabilities in terms of how information has been used. Moreover, the Google Play approach notifies users of changes only when a new version of the app is pushed into the store. In our approach, such changes are determined using the provenance record which is independent from new versions of applications, devices or services (e.g. change in the server infrastructure with regards to manipulation of user's data triggers notification to user).

4 Demonstration Content

In the demonstration we will illustrate the behaviour of the system using the two scenarios described above. In the Bus Stop scenario we will highlight the capabilities of the NFC device based on prospective provenance. A short presentation video of this scenario can be viewed at our Trusted Tiny Things website[10]. In the car blackbox scenario we will demonstrate how retrospective provenance is used to infer the capabilities of the telemetry box. Finally, we will showcase our notification service by changing the way that the insurance service operates (it will begin to share sensor data with car manufacturers). We will demonstrate how our system would detect the change and infer new capabilities associated with this change (i.e. confidential data is now shared with a third-party company).

References

1. Weber, R.H., Weber, R.: Internet of Things. Springer, New York (2010)
2. Bizer, C., Heath, T., Berners-Lee, T.: Linked data-the story so far. Int. J. Semant. Web Inform. Syst. **5**(3), 1–22 (2009)
3. Lebo, T., Sahoo, S., McGuinness, D., Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., Soiland-Reyes, S., Zednik, S., Zhao, J.: Prov-o: The prov ontology. W3C Recommendation, 30 April 2013

[10] http://t3.abdn.ac.uk.