

Exposing Internet of Things Devices via REST and Linked Data Interfaces

Tobias Käfer, Sebastian Richard Bader, Lars Heling, Raphael Manke, and
Andreas Harth

Institute AIFB, Karlsruhe Institute of Technology (KIT), Germany
{tobias.kaefer|sebastian.bader|harth}@kit.edu,
{lars.heling|raphael.manke}@student.kit.edu

Abstract. Web technologies are widely used open standards to interconnect devices on virtually any platform. The W3C's Web of Things effort wants to bring web technologies to the Internet of Things to address interoperability challenges on the Internet of Things. In this paper, we report on four independently developed implementations whose aim was to expose Internet of Things devices via web technologies, more concretely REST and Linked Data interfaces. While they were all developed with the same goal in mind, and with technologies that promote uniformity on the interfaces, the implementations still exhibit different schemas and architectures with different communication patterns.

1 Introduction

In the last couple of years, developments of embedded devices have made networked sensors and actuators from various vendors in various domains affordable. The adoption and the deployment of such devices is considerable. To provide a common term for such networked sensors and actuators, the term Internet of Things (IoT) has been coined. To expose the IoT device data and functionality on a network interface, the different vendors in different domains came up with different ways of describing the data and interacting with the devices. This heterogeneity makes an integrated processing of data from various devices a hard task: For example, manufacturing machine suppliers equip their machines with numerous networked sensors and actuators. While the machines from the individual suppliers can be optimised, holistic optimisation over the whole shop floor of a manufacturer is hindered because of interoperability problems. Similarly, in smart home scenarios, there are solutions like ZigBee Light Link¹ (Philips Hue², Osram Lightify³), which give interoperability on a local scale in the light domain. But this heterogeneity is a challenge when building applications that make use of devices from different domains and vendors. This challenge needs

¹ <http://www.zigbee.org/zigbee-for-developers/applicationstandards/zigbee-light-link/>

² <http://www.meethue.com/>

³ http://led.osram.de/led_de/lightify/lightify-produkte/index.jsp

to get addressed, at latest for the Industrial Internet of Things (also known as Industrie 4.0), where connections between heterogeneous systems of different companies have to be built, which are driven first and foremost by business partnerships and not common technology.

Before applications can be built that make use of IoT devices, the devices' data has to be made accessible. In making the data accessible, one faces two degrees of freedom:

1. How to interact with the device to obtain the data?
2. How to describe and represent the data?

We propose to use established Web technologies to build IoT systems, thus embracing the W3C's Web of Things effort (see also Section 2). Specifically, we use HTTP (the Hypertext Transfer Protocol), an implementation of the REST (representational state transfer) architectural style, for the interaction and RDF (the Resource Description Framework), a data model for representing the data. This combination is also called Linked Data. Although we take those decisions upfront, the technologies HTTP and RDF still present us with choices for both degrees of freedom even for 1.

Before the advent of the IoT, the proliferation of mobile devices gave rise to cloud applications that expose RESTful APIs (Application Programming Interface) for mobile apps, which mainly consume data. The architecture of such cloud applications are characterised by a central server, on which all users' data is maintained. Such centralisation of data brought privacy issues and led to data silos with a proprietary data model, where it is hard to get data out. This architecture with a central server is also a way to connect IoT devices. While in the mobile apps scenario, the distributed and power-limited device is mainly the consumer of data, on the IoT, the opposite is true: We have highly distributed, unreliable and power-limited devices to whose data access is to be organised.

In this paper, we want to present two approaches to get data and functionality from sensors and actuators accessible using web technologies such they can get integrated in applications. In our case, those applications are research prototypes that consume data from Linked Data interfaces. The paper is a report on four independently developed implementations, two for each approach, that we carried out from scratch to provide Linked Data access to our IoT devices. The paper is structured as follows: In Section 2, we give basic definitions and introductions to the technical terms we use in the remainder of the paper. In Section 3, we describe four implementations to get sensors and actuators into the Internet of Things. Last (Section 4), we summarise and conclude.

2 Preliminaries

The Internet of Things is under active development, with many standardisation organisations (e.g. the ISO/IEC Internet of Things working group) and industry-led consortia (e.g. the AllSeen Alliance⁴, the IPSO Alliance⁵, the Open

⁴ <http://allseenalliance.org/>

⁵ <http://www.ipso-alliance.org/>

Interconnect Consortium⁶, the Open Connectivity Foundation⁷, combining the two former organisations, and the Industrial Internet Consortium⁸) that work on creating and establishing dedicated IoT standards. Given that these efforts are rather young and their standardisation efforts in constant flux, we do not attempt to provide a survey of the different approaches. Rather, we embrace the W3C’s effort around the “Web of Things”⁹, assuming the Internet of Things is going to be using established web technologies such as URIs and HTTP. We therefore share the aim and the base technologies with the Web of Things effort, despite our research focus is different from the standardisation work of the Web of Things group. While the Web of Things effort is concerned a lot with descriptions of offerings and capabilities of things that *may* have Linked Data interfaces, we assume Linked Data interfaces and are concerned with the semantic data provided by the thing, the interaction with the thing and the execution of programs that employ the thing. On the other hand, we note that there are other less widely deployed web standards such as WebSockets, which can also be used to connect IoT devices [11], which are also recognised by the Web of Things effort.

The web can be described as a system that implements the architectural style of Representational State Transfer (REST) [9]. We use the terminology introduced in [9] to describe and contrast the different approaches in this paper for accessing IoT devices as part of applications. Broadly speaking, applications following the RESTful architectural style consist of data, connectors and components. In the remainder of the section, we address each in turn, starting with data.

In RESTful architectures, the central notion is that of a “resource”; a resource is any identifiable and referenceable thing in the context of an application. To reference resources, they need names. On the web, Uniform Resource Identifiers (URIs) [1] serve as names for resources. In our examples, we use URI templates as specified in [10]. A resource has a state that may change over time. To access and transfer the state of resources between components, we need a way to represent this state. Those resource representations need to include references to resources to allow for linking and discovering previously unknown resources in decentralised networked environments such as the web or the Web of Things.

All of the presented approaches use the Resource Description Framework¹⁰ (RDF) to represent the state of resources. RDF consists of subject-predicate-object triples. A set of such triples forms a RDF graph. A RDF document is a document that encodes a RDF graph. In a triple, the subject is a resource identifier, which can be a URI, or a blank node, which is a document-scoped local identifier. The URI-identified predicate defines the type of the relation between the subject and the object. In object position of a triple, there can be

⁶ <http://openinterconnect.org/>

⁷ <http://openconnectivity.org/>

⁸ <http://www.iiconsortium.org/>

⁹ <http://www.w3.org/Submission/wot-model/>

¹⁰ <http://www.w3.org/TR/rdf11-concepts/>

either a URI, or a blank node, or a Literal, which are used for values such as strings or numbers.

Connectors are concerned with the transfer of representations of state of resources between components [9]. Connectors provide, in other words, the means for communication. REST defines the following connectors: client, server, cache, resolver, tunnel. From this list, only the two first are relevant in the scope of this paper.

We use the Hypertext Transfer Protocol (HTTP) [8] as our communication protocol. An interaction between components in HTTP consist of a request-response pair of messages. A HTTP client connector starts the communication by sending a request to a specified server connector which in turn answers with the response message. Depending on the type of request, the message includes a body. HTTP offers several different types of requests, distinguished by their method: Most common are GET for retrieving a resource state representation, PUT for creating a resource or overwriting its state representation and DELETE to delete a resource. Then, there is POST, which can be used to append to a collection of resources, or to invoke a data processing task on the server. Therefore, to provide an interface to an actuator, the PUT request can be used in the following fashion: if we are to allow change to the state of an actuator using a PUT request, then the payload of the PUT request contains a description of the desired resource state. This description of the desired state can be derived by first retrieving a description of the current state using a GET request, and then modifying this description to reflect the desired state.

In REST, a connector to send or answer requests resides on a so-called component. Among intermediaries (see Section 3.2), there are user agents (the client program initiating a request to a resource) and origin servers (the source of authoritative information on the resource). The terms client and server are only concerned with the roles in an exchange of one request-response pair, as one component may act as client in one exchange and as server in a different exchange. The decision of the components mainly acts as server and which mainly act as clients, is for the designers of the application to decide. For our implementations, we started out from scratch and different circumstances made us go different routes, which we describe in Section 3.

The Linked Data principles [2] are a set of practices to publish data on the web, which make use of the combination of RDF and HTTP: To allow for data retrieval of resources, the first two principles suggest to use HTTP URIs to identify things. The third principle recommends to use standard-conforming data representations such as RDF. To allow for the discovery of new information, the fourth principle advocates to include links in the representation to other relevant data.

While the Linked Data principles are about data publishing, the interaction with web resources that are described in RDF was not in their scope. The W3C's Linked Data Platform (LDP)¹¹ specification closes the gap between the HTTP and the RDF specification for the RESTful interaction with web resources. The

¹¹ <http://www.w3.org/TR/ldp/>

recommendation defines Linked Data Platform Resources and Linked Data Platform Containers. A LDP Resource guarantees a minimal set of common read operations and specifies how servers publishing such resources need to react to HTTP requests targeting the resource. LDP containers are collection LDP Resources with additional functionalities for creating new resources.

3 Approaches and Implementations

We implemented two different approaches: The first two implementations directly connect the devices with sensors/actuators to the web. The second two implementations use an intermediary. All approaches expose the information in the form of web resources. We describe the approaches and implementations in this section. We cover in detail the involved components and their interaction. We also name the vocabularies that are used for describing the data for each implementation, but omit the detailed presentation of the RDF data, as this is not in the focus of the paper. Yet, we provide links to the source code of each implementation for the interested reader.

3.1 Direct Access to the Device with the Sensor/Actuator

In this section, we present two implementations that implement the REST component type “origin server”, i. e. there is direct access to the definitive source of information about the resources. The origin server thus implements a REST “server connector”.

Connecting two Modules of a Tessel 2 In this section, we describe how we built a REST and Linked Data interface for the Tessel2¹². The corresponding code can be found online¹³. Tessel2 is a PCB (Printed Circuit Board) that has been developed to easily build networked sensor/actuator devices. The board comes with USB, Ethernet and Wi-Fi connection and two sockets for proprietary sensor and actuator modules. The PCB runs Node.js¹⁴ such that one can program the PCB in JavaScript. The modules come with corresponding libraries. We use a Tessel2 PCB with the ambient module, which includes sensors for light and sound, and the relay module, which includes two relays to switch 240V at 5A, see Figure 1.


On the HTTP interface, we describe the board with its two sensors using the LDP vocabulary. The root resource thus points to two resources describing the modules using the `ldp:contains` property. The ambient module points to the resources for the sensors in the same manner. The sensors are described as observation from the data cube vocabulary¹⁵. The sensor readings are connected to the sensor using a custom property. The relay module also uses the

¹² <http://tessel.io/>

¹³ <http://github.com/kaefer3000/t2-rest-relay-ambient>

¹⁴ <http://nodejs.org/>

¹⁵ <http://purl.org/linked-data/cube#>

Fig. 1. Component Diagram of the Implementation using the Tessel 2.

`ldp:contains` property to point to its relays. The relay’s state (on/off) is also connected to the relay itself using a custom property. Using a PUT request on the relay’s URI, the relay can be switched on and off.

Table 1. API Description of the Implementation for the Tessel 2.

URI Template	Method	Description
/	GET	Returns an overview over the available modules.
/ambient/	GET	Returns an overview over the available sensors of the ambient module.
/ambient/{sensor}	GET	Requests a reading from the sensor, i.e. light or sound, and returns it.
/relay/	GET	Returns an overview over the available power switches.
/relay/{n}	GET	Returns the status for the n -th power switch.
/relay/{n}	PUT	Sets the status for the n -th power switch.

We implemented the HTTP interface using the Express¹⁶ for building Web APIs in JavaScript. We use the RDF serialisation JSON-LD as way to represent the data.

Connecting a SenseHat In this section, we describe how we built a REST and Linked Data interface to interact with a Sense Hat¹⁷, an add-on board for the Raspberry Pi¹⁸, a single-PCB computer. The Sense Hat board offers several sensors (humidity, gyroscope, accelerometer, magnetometer, temperature, barometric pressure), and a programmable LED matrix. The code can be found online¹⁹.

To describe the data on the interface, we use parts of the popular SSN ontology. The SSN (Semantic Sensor Network) ontology from the W3C’s Semantic

¹⁶ <http://expressjs.com/>

¹⁷ <http://www.raspberrypi.org/products/sense-hat/>

¹⁸ <http://www.raspberrypi.org/>

¹⁹ <http://git.scc.kit.edu/ujdpo/sensehat>

Sensor Network Incubator Group²⁰ proved not to be specific enough. The SSN ontology describes the way several devices with several sensors work together. A way to specifically describe one device is not part of the SSN ontology, nor how the data of a sensing device can be structured. Therefore, we built a vocabulary ourselves by extending the SSN ontology by the class `LedMatrix`. Each sensor on the Sense HAT board is an instance of the SSN class `SensingDevice`, and the LED matrix of `LedMatrix`. A `SensingDevice` is connected to the quantity kind (such as humidity and temperature) it observes using the `featureOfInterest` property. On top of that a `SensingDevice` is connected to the RDF literal representing the measured value using the `observationValue` property and to the URI representing the corresponding unit using the `observationUnit` property. Those descriptions can make use e.g. of `dbpedia`²¹ or the `QUDT` ontology²². For an instances of the `LedMatrix` class, the number of LEDs (instances of the class `SingleLed`) in X and Y direction can be given. Using the `hasLed` property, the connection between a `SingleLed` and a `LedMatrix` can be stated. For each `SingleLed`, the X and Y coordinate can be given using the `x-coordinate` property and the respective property for the Y direction. The colour of each `SingleLed` can be defined using three properties, which define the values in a range between 0 and 255. The property for red is `red-color` and there are analogous properties for blue and green. We serve dereferencable URLs for the data about the Sense Hat as described in Table 2.

Table 2. API Description of the Implementation for the SenseHat.

URI Template	Method	Description
/	GET	Returns information on the Raspberry Pi linking to the installed boards, e.g. the Sense Hat.
/sensehat/	GET	Returns information on the Sense Hat linking to the sensors and the LED matrix.
/sensehat/led/	GET	Returns information on the LED matrix linking to the individual LEDs. Additionally, it returns all LEDs' current value.
/sensehat/led/	PUT	Overwrites the information on the LED matrix. Can be used to set the values of the LEDs.
/sensehat/led/{x}/{y}	GET	Returns the colour of the LED at (x, y) on the LED matrix.
/sensehat/{sensor}	GET	Returns the value of a sensor.

We implemented the interface to the Sense Hat of the Raspberry Pi in Python. The Sense HAT python library²³ allowed to program access to the

²⁰ <http://www.w3.org/2005/Incubator/ssn/>


²¹ <http://dbpedia.org/>

²² <http://www.qudt.org/>

²³ <http://pythonhosted.org/sense-hat/>

sensors of the Sense Hat and the LED matrix. We used the framework Flask²⁴ to implement the interaction with the resources. The library RDFlib²⁵ served in the serialisation and deserialisation of RDF data. Upon a request to a URI on the REST interface, the corresponding data is retrieved from the Sense Hat, described in RDF and then sent as a response.

Fig. 2. Component Diagram of the Implementation using the SenseHat.



Some extra focus must be taken on the way the LEDs can be changed as several LEDs are addressed in the same request. But as the single LEDs do not have an unique identifier, the only available method is to specify them by their x and y coordinate of the matrix. Two triples, one for the x-axis and one for the y-axis, define each LED respectively. On subject position of the triples one can use blank nodes. The following triples contain the thereby specified blank nodes as subjects and determine the new behaviour of the corresponding LED.

3.2 Access via an Intermediary to the Device with the Sensor/Actuator

In addition to the direct access to the information source we present two implementations that use an intermediary: There are two components: one component that has a sensor/actuator connected (we call the sensor/actuator-bearer) and a client connector. Then, there is another component, an intermediary, with a server connector exposing a resource that represents the state of the sensor/actuator connected to the other component. The state of this resource is periodically updated by the first connector using PUT requests.

What is this Approach in REST? In REST terms, this pattern is hard to grasp. One could argue that there is a resource on the first connector that cannot get accessed directly, because the sensor/actuator-bearer has no server connector, and that there is a logical correspondence between the not directly accessible resource and the accessible resource on the second component. Next, we discuss the origin server and the different intermediary component type REST offers and why they do not fit the pattern:

²⁴ <http://flask.pocoo.org/>

²⁵ <http://github.com/RDFLib>

Origin Server An origin server is defined as “the program that can originate authoritative responses for a given target resource” [8]. While our intermediary is the source for information on a given resource, the authoritative source is the sensor/actuator-bearer. Moreover, the term origin server would fit for the latter, because a characteristic of an origin server is to “be the ultimate recipient of any request” [9]. But the sensor-bearer cannot be called origin server, because it does not have a server connector, therefore, it cannot give authoritative *responses*.

Proxy A proxy is defined as “a message-forwarding agent that is selected by the client [...] to receive requests [...] and attempt to satisfy those requests via translation through the HTTP interface” [8]. While our intermediary is not selected by the client, because the client does not know that the source of the information is not the intermediary, our scenario looks like the example proxy in [9, Fig. 5-10 c].

Gateway A gateway is defined as “intermediary that acts as an origin server for the outbound connection but translates received requests and forwards them inbound to another server or servers” [8]. While in our case the intermediary indeed acts as an origin server, it does not forward the request. On the other hand, the encapsulation of other services is one of the examples of the gateway, and it may communicate with other servers using any protocol [8]. Thus, the gateway could be regarded as the closest fit.

The Intermediary Strategy as a Way of Addressing IP-layer Issues

The approach with a central server is necessary e. g. in situations where the ongoing transition from IPv4 [13] to IPv6 [5] is solved using so-called DS-lite [7] connections: In IPv4, every consumer gets one single an IPv4 address from his Internet Service Provider (ISP), which is reachable from the Internet. For a customer to nevertheless connect multiple devices to the internet, network address translation (NAT) [14] is used: The customer employs a router, which gets the one IPv4 from the ISP. The devices on the local network can connect to the Internet by sending a request to the router, which replaces the local IP that initiates the connection with his own IP, and assigns a free port to the connection. The router then forwards data that comes to this port from the Internet to the local IP and port that initiated the connection. All computers in the local network thus appear as one computer on the Internet. Connections initiated from the Internet can only reach a local device if the router is configured to do port forwarding: The router maintains a mapping from his own ports to a tuple of local IP and port of local devices and forwards traffic accordingly.

As we go from IPv4 to IPv6, many ISPs employ a technique called Dual Stack lite (DS-lite) [7]: Instead of giving both an IPv4 and IPv6 address to their customers (Dual Stack [12]), ISPs only give IPv6 addresses to their customers. If the customers request a connection to an IPv4 address on the Internet, the ISPs tunnel the traffic accordingly. Conversely, this means that there cannot be a IPv4 connection initiated from the Internet, because the customer does not have an IPv4 address. This is particularly an issue for devices on cellular network,

where typically only IPv4 is deployed. Moreover, if no further port forwarding is employed, only connections to IPv6-enabled local devices can be initiated from the Internet. However, not only the deployment of DS-lite is an argument for the proposed architecture, but also the fact that it alleviates the necessity to have access to the router with the necessary rights to configure the port forwarding.

Addressing Device Limitations using the Intermediary Strategy We imagine two scenarios here:

- The continuous availability of the device is not guaranteed, but a high availability of the data is important. An off-the shelf caching intermediary for HTTP would be sufficient here, though
- The intermediary can answer more complex requests on top of HTTP based on the same data. To have the data both on the device and the component, which is proposed to be an intermediary, would be an unnecessary duplication of data.

Connecting Weather Sensors and a 433 MHz Transceiver In this section, we describe how we built an REST and Linked Data interface to (a) the functionality a 433 MHz transceiver provides (b) a directly connected temperature sensor. The code can be found online²⁶. In our scenario, the transceiver wirelessly controls two power sockets, and receives data from a weather station, more concretely, temperature and humidity values. The transceiver is connected to the GPIO pins of a Raspberry Pi. Moreover, there is another temperature and humidity sensor directly connected to the GPIO pins of the Raspberry Pi. In this implementation, the access to the client is implemented indirectly: The Raspberry Pi exchanges data with a central server, which stores data about the current state of the sensors and sockets. For the sockets, the data on the server can also mean the desired state. One rationale for this approach is to alleviate the necessity to directly access the clients to retrieve sensor data and control actuators.

The implementation describes a sensor reading as **Observation** from the SSN ontology. To describe the value, we use Wikidata²⁷ and the Smart Appliances Reference ontology²⁸. The location of the reading is described using the WGS84²⁹ ontology. The values are described using XSD data types³⁰.

The client periodically takes snapshots from the sensors directly connected to the Raspberry Pi. Additionally, a 433Mhz transceiver is used to first, wirelessly receive temperature and humidity values from a weather station sensor, and second, wirelessly control two power sockets. The weather station sends the current values for temperature and humidity every five minutes.

²⁶ <http://github.com/Lars-H/home-automation>

²⁷ <http://www.wikidata.org/>

²⁸ <http://ontology.tno.nl/saref.ttl>

²⁹ http://www.w3.org/2003/01/geo/wgs84_pos

³⁰ <http://www.w3.org/TR/xmlschema-2/>


All data from the sensors is enriched with meta data, sent to a server where they get stored in a relational database. Moreover, there are records in the data base for the sockets. The relational database is made accessible using a RESTful API, which allows for requests to obtain and to store data. On the interface the sensors and actuators are addressed using URIs. Sensor data can be requested with a HTTP GET request on the corresponding URI and action for actuators can be sent using HTTP POST request to the corresponding URI.

To implement the corresponding tasks, there are three parallel processes running on the Raspberry, two processes to handle the sensors and one for handling the actuators:

- Process 1 periodically requests the current value for temperature and humidity from the connected sensor. It enriches the sensor data with meta data and sends the data to the database server.
- Process 2 is the (event-triggered) observer for the wireless connected sensor. Any time a new data record is received from the sensor, it is also enriched with meta data and sent to the server.
- Process 3 periodically requests the current status from the database for the actuators. In case the status changes, it will perform the corresponding task and update the status on the database.

The database is a MySQL database which can be accessed using an HTTP interface to store and retrieve data. The REST-interface is realised as a server using the Flask Framework³¹ for building RESTful APIs in Python. A description of the API can be found in Table 3. The API supports content negotiation for different RDF serialisations.

Fig. 3. Component Diagram of the Implementation for the Weather Sensors and the 433 MHz Transceiver.



Connecting a SensorTag The goal of this approach is to publish various outputs of a Texas Instruments SensorTag CC2650³². The SensorTag delivers sensors for temperature, humidity, acceleration, and light which can be accessed via a Bluetooth Low Energy interface. The range of the bluetooth connection limits the distance between the sensor itself and an according control unit where

³¹ <http://flask.pocoo.org>

³² <http://www.ti.com/sensortag>

Table 3. API Description of the Implementation for the Weather Sensors and the 433 MHz Transceiver.

URI Template	Method Description	
<code>/sensor/outside/{type}</code>	GET	Returns the latest record of the sensor value of type humidity or temperature for the wirelessly connected weather station.
<code>/sensor/inside/{type}</code>	GET	Returns the latest record of the sensor value of type humidity or temperature for the directly connected sensors.
<code>/actuator/plug{n}/status</code>	GET	Get the current status for the n -th wireless power plug.
<code>/actuator/plug{n}/status</code>	PUT	Set the status for the n -th wireless power plug.

at the same time the positioning of the SensorTag should not depend on space requirements of a powerful hardware. Further on, we want to query the real-time observations while also be able to access the whole data for analytical tasks on the historical data.

We use the **Observation** from the SSN ontology [4] to describe a sensor reading, and describe the data using XML Schema datatypes, vCard³³, and QUDT. We connect the SensorTag to a Raspberry Pi Model B equipped with a Bluetooth 4.0 dongle. The Raspberry Pi is programmed to periodically request data from the SensorTag using bluepy³⁴. Bluepy supplies JSON data. Another process (“Administration Shell” in Industrie 4.0 terminology [6]) checks whether this data from the SensorTag has changed and if so, the process lifts this sensor data to RDF and represents it as an observation according to the SSN ontology [4], XML Schema datatypes, and vCard. This administration shell sends the RDF data to a collection resource on a Linked Data Platform implementation, more precisely to a `ldp:BasicContainer`. We use Apache Marmotta as an implementation of the Linked Data Platform, as it both offers a Linked Data interface for RESTful interaction and a SPARQL interface for complex queries demanded by our analytics use-case. Especially for the requirements of a profound analytical application, we assume the Raspberry Pi not powerful enough to store the amount of data and do a sufficient query processing. That’s the reason why we have the Apache Marmotta instance running on an Ubuntu 14.04 based virtual machine, separating the data management from the data producer. The code can be found online³⁵.

4 Discussion and Conclusion

We described four independently developed implementations of a REST and Linked Data interface for IoT Devices. The interface is thus built on open stan-

³³ <http://www.w3.org/TR/vcard-rdf/>

³⁴ <http://github.com/IanHarvey/bluepy>

³⁵ http://github.com/sebbader/KSRI-KM_STEP

Fig. 4. Component Diagram of the Implementation using the SensorTag.




Table 4. The Linked Data Platform API as implemented by Apache Marmotta.

URI Template	Method	Description
/marmotta	GET	Returns information and configuration options of the Marmotta Platform
/marmotta/ldp	GET	Returns information about the root container
/marmotta/ldp/ldbc	GET	Responds with information on the used Linked Data Platform Basic Container which serves as the central access point in this implementation
/marmotta/ldp/ldbc	POST	Creates a new Linked Data resource
/marmotta/ldp/ldbc/{x}	GET	Returns information about the data object x
/marmotta/ldp/ldbc/{x}	PUT	Overwrites an existing resource or creates a new one

dards. We hid the technology fragmentation of the sensors and actuators behind a uniform interface and data model. Despite this uniformity as aim of all implementations, we ended up with four implementations varying in vocabularies and interaction direction with the device that bears the sensor/actuator.

In terms of interaction direction, the implementations can be categorised according to whether there is direct access to the device with the sensor attached or not. While the direct access would be the expected pattern in a RESTful architecture, there are strong reasons why in some settings the access via an intermediary is reasonable. The intermediary again provides a uniform interaction direction in all implementations.

The different implementations resulted in vocabulary heterogeneity, although the use-cases are very similar. This is because while there are elaborate established vocabularies to describe sensors and values, there are no established simple constructs to describe properties of physical objects. Part of the W3C’s effort of the Thing Descriptions³⁶ is to come up with such constructs. They note that already on the capability description level, ontologies do not much agree on terminology [3]. In the meantime, the use of the semantic data model RDF allows for employing reasoning techniques in the applications that interact with the devices to integrate the data described using different vocabularies.

³⁶ http://www.w3.org/WoT/IG/wiki/Thing_Description

Acknowledgements

This work is partially supported by AFAP, a BMBF Software Campus project (FKZ 01IS12051), the BMBF ARVIDA project (FKZ 01IM13001G), and the BMWi project STEP (FKZ 01MD16015B).

References

1. Berners-Lee, T, Fielding, R, and Masinter, L: *Uniform Resource Identifier (URI): Generic Syntax*. RFC 3986 (INTERNET STANDARD). Updated by RFCs 6874, 7320. Internet Engineering Task Force (2005).
2. Berners-Lee, T: Design Issues – Linked Data, (2009). <http://www.w3.org/DesignIssues/> (visited on 09/16/2016)
3. Charpenay, V, Käbisch, S, and Kosch, H: Introducing Thing Descriptions and Interactions: An Ontology for the Web of Things. In: Proceedings of the 1st Workshop on SemanticWeb technologies for the Internet of Things (SWIT) at the 15th International Semantic Web Conference (ISWC) (2016). in print
4. Compton, M: The SSN ontology of the W3C semantic sensor network incubator group. Web Semantics: Science, Services and Agents on the World Wide Web 17 (2012)
5. Deering, S and Hinden, R: *Internet Protocol, Version 6 (IPv6) Specification*. RFC 2460 (Draft Standard). Updated by RFCs 5095, 5722, 5871, 6437, 6564, 6935, 6946, 7045, 7112. Internet Engineering Task Force (1998).
6. Dorst, W, ed.: Umsetzungsstrategie Industrie 4.0. Ergebnisbericht der Plattform Industrie 4.0, Plattform Industrie 4.0. (2015). http://www.bitkom.org/Bitkom/Publikationen/Publikation_5575.html
7. Durand, A, Droms, R, Woodyatt, J, and Lee, Y: *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*. RFC 6333 (Proposed Standard). Updated by RFC 7335. Internet Engineering Task Force (2011).
8. Fielding, R and Reschke, J: *Hypertext Transfer Protocol (HTTP/1.1): Message Syntax and Routing*. RFC 7230 (Proposed Standard). Internet Engineering Task Force (2014).
9. Fielding, RT: Architectural styles and the design of network-based software architectures. Diss. University of California, Irvine (2000)
10. Gregorio, J, Fielding, R, Hadley, M, Nottingham, M, and Orchard, D: *URI Template*. RFC 6570 (Proposed Standard). Internet Engineering Task Force (2012).
11. Merkle, N, Kämpgen, B, and Zander, S: Self-Service Ambient Intelligence Using Web of Things Technologies. In: Proceedings of the 1st Workshop on Semantic Web Technologies for Mobile and Pervasive Environments (SEMPER) at the 13th Extended Semantic Web Conference (ESWC) (2016)
12. Nordmark, E and Gilligan, R: *Basic Transition Mechanisms for IPv6 Hosts and Routers*. RFC 4213 (Proposed Standard). Internet Engineering Task Force (2005).
13. Postel, J: *Internet Protocol*. RFC 791 (INTERNET STANDARD). Updated by RFCs 1349, 2474, 6864. Internet Engineering Task Force (1981).
14. Srisuresh, P and Egevang, K: *Traditional IP Network Address Translator (Traditional NAT)*. RFC 3022 (Informational). Internet Engineering Task Force (2001).