Describing the Internet of Things with an Ontology: The SusCity Project Case Study

David Perez Abreu*, Karima Velasquez*, Alexandre Miguel Pinto†, Marilia Curado* and Edmundo Monteiro*
*Department of Informatics Engineering
University of Coimbra, Portugal
Email: {dabreu,kcastro,marilia,edmundo}@dei.uc.pt
†Department of Informatics
University of Lisbon, Portugal
Email: amdpinto@ciencias.ulisboa.pt

***Abstract*—The Internet of Things comprises a network of physical objects, like sensors and actuators, collecting and exchanging data. Given the importance of the information exchanged in these environments, the communication infrastructure becomes a critical point that needs to be managed optimally, while providing high-performance levels to end users. To guarantee the correct interaction between the different procedures intended for the optimization of the communication infrastructure, a standard and flexible representation of the data related to the network is necessary. Many of the services could be monitored through the web, thus using standard Web languages with a rich expressive power, such as the languages used in the Semantic Web, would allow for the reification of interoperable descriptions.**

This is one of the main challenges faced by the SusCity project, which aims at the realization of the Smart City paradigm in Lisbon, Portugal, by designing a dashboard to monitor smart building, transportation, and smart grids. This paper presents an ontology for the Internet of Things infrastructure tailored to the needs of the SusCity project. Furthermore, different kinds of evaluations were performed to corroborate the correctness of this ontology, including potential infrastructure optimization objectives like low latency and high resilience.

I. Introduction

Society is currently undergoing a shift in the way daily activities take place. Automation for common activities is leading to the instantiation of the Smart City paradigm [1] [2], where services are made more efficient by combining different solutions such as Information and Communication Technologies (ICT) and the Internet of Things (IoT) [3] [4]. Several efforts are being made in this direction, to help in the transformation of regular cities into Smart Cities. The SusCity project [5] is one of these examples, with the main goal of developing an operating model for the city of Lisbon, Portugal; and also from its resources (e.g. energy, materials) consumption, in order to plan effective strategies for resources usage.

In this scenario, data from different sources can be collected in an automated fashion from a wide range of devices (e.g. smartmeters, smartphones, sensors, data loggers) from various manufacturers. This data can be of different natures (e.g., real-time or not, aggregated or single) which imposes new challenges for the ICT platform, like dealing with the heterogeneity of the described environment providing a uniform language to enable the communication among the devices (i.e. communication protocol); furthermore, it is important to define, in a standard fashion, the communication infrastructure itself, which will allow to use different mechanisms to manage and optimize its performance.

To manage the communication infrastructure, it is important to have a detailed description of the network that allows the aforementioned mechanisms to make the necessary adjustments to obtain the desired optimization (e.g. lower latency, lower energy consumption, higher resilience). The heterogeneous nature of the network and devices forces the use of a standard that can guarantee the correct interaction among the different mechanisms as well as a proper update process for the description model itself.

Using a standard mechanism to describe the communication infrastructure becomes evident in this scenario as a way to deal with its heterogeneity. Moreover, this description must be shared among different applications that usually communicate using web services. One possible solution is using a well-known standard language, such as Resource Description Framework Schema (RDFS) [6], Web Ontology Language (OWL) [7], or Terse RDF Triple Language (Turtle) [8] to share the description of the communication infrastructure between the different applications previously described. This would facilitate the usage of data input to intelligent protocols, services, and user applications, enhancing their operation.

In this paper, an ontology for the Internet of Things infrastructure in Smart City scenarios is proposed; in order to unify the access to the data, making it possible to collect, in a standard and easy way, the information that protocols, services, and user applications need to make smart decisions to optimize the performance of the network, according to the requirements inherent to Smart Cities.

The proposed ontology was successfully tested against consistency issues and also populating and querying it using SPARQL Protocol and RDF Query Language (SPARQL) [9]. Particularly, two optimization scenarios (lower latency and high resilience) were used in the design of the example queries as a way to illustrate the potential impact of using this type of solution to describe a Smart City environment to optimize the performance of the infrastructure.

II. Related Work

In order to deal with the complexity of the environment previously described, different research proposals have been considered. One of these approaches is using semantic technologies, particularly ontologies, to deal with the heterogeneity of the actors involved in this environment. This section presents a review of previously conducted researches that apply ontologies in the field of IoT.

Two complete and well-detailed ontologies that model services in the IoT are presented by De et. al. [10] and Byun et. al. [11]. The first ontology describes general services whereas the second one is focused on social network services. Although the ontologies exposed in these works describe in detail the services, they lack the description of the devices and the infrastructure of the IoT, two important aspects that are covered in this paper.

Manate et. al. [12] present a survey of different approaches to using semantic technologies in the IoT. In particular, the authors take into consideration a subset of studies focused in ontologies applied to the IoT to support their ideas. Also, the possibility of using ontologies and multi-agents systems in the context of IoT and Smart Cities is mentioned. As stated before, the work of Manate et. al. constitutes a survey whilst this work presents a concrete case study of applying semantic technologies to describe the IoT.

Hachem et. al. [13] discuss how to use semantic technologies as the main actor of middleware to address the diversity and scalability of objects in IoT. Additionally, they propose two ontologies, the first one describes devices and their functionalities in the IoT, while the second one focuses on modeling the domain of the physical (real) environment in the IoT; specifically, this ontology tries to model information about real world physical concepts and their relations. The main difference between the work of Hachem et. al. and this research is based on the fact that the first one is focused on the implementation of middleware, for which they use too tailored ontologies, that can not be applied to generic Smart City scenarios.

Wang et. al. [14] present an ontology for the IoT domain and also discuss how their ontology can be used to support tasks like service discovery. In their ontology, they classify the services in Simple Object Access Protocol (SOAP) and Representational State Transfer (REST) service types and also include a description of the network components and the Quality of Service (QoS) provided by them. This ontology is more focused on the services that can be provided in the IoT instead of in the infrastructure itself.

Later, in another research in this field, Wang et. al. [15] introduce an ontology to modularize physical objects in the digital world. Their ontology describes the common objects deployed in a home environment (i.e. appliances and furniture) and some services that use the described object. Despite that the previous ontology is well-defined, it only takes into consideration the actors involved in a smart-home.

As stated before, the works reviewed in this section do not provide a thorough description of an environment that resembles the demands of the SusCity project, specifically from the infrastructure point of view. Thus, a new ontology must be developed following the interests of Smart City scenarios in order to fully characterize the physical environment and eventually use this description on different processes regarding the infrastructure; all this taking into consideration that it will have an immediate impact on the user applications that will influence the quality of life of the citizens, financial aspects for business, and government activities.

III. An Ontology for the IoT Infrastructure

This section contains the characterization of the entities included in the ontology as well as their relationships. The main objective of this proposal is to build a complete and detailed ontology to model all the components involved in the IoT infrastructure that satisfies the requirements of the SusCity project. Particularly, there is particular emphasis on the communication devices with the goal of gathering information that enables the decision-making process of different mechanisms that guarantee the proper function of the infrastructure. A complete version of the ontology discussed in this section is available online [16].

The IoT infrastructure is represented by an entity (*IoT_Infrastructure*) that includes different elements attached to it. After evaluating the interests and needs of the SusCity project, four additional entities were identified (*Device*, *Link*, *Interface*, and *Metric*). Two additional classes (*Action* and *Location*) were added to simplify the modeling of some required information that is going to be detailed later in the paper.

The main class hierarchy is shown in Figure 1. Several facts are of interest at this point; for instance, the support for both wired (e.g. xDSL, Fiber, Ethernet) and wireless (e.g. Bluetooth, Cellular, NFC, WiFi, WiMAX, RFID, ZigBee) communication interfaces and links, the wide variety of heterogeneous devices included in the infrastructure (e.g. Communication, Location, Multimedia, Transducer), and also the inclusion of both devices and link metrics that will enable the monitoring of the network status.

Once the class hierarchy is described, it is important to detail the classes in the ontology further. Since the main goal of this work is to design and implement mechanisms to manage and optimize the function of the communication infrastructure, this becomes the main component of the developed ontology. The description of the *IoT Infrastructure* class is provided below.

$$\begin{aligned}
&\text{IoT_Infrastructure} \sqsubseteq \text{Thing} \\
&\text{IoT_Infrastructure} \sqsubseteq \exists\ \text{isComposedBy Interface} \\
&\text{IoT_Infrastructure} \sqsubseteq \exists\ \text{isComposedBy Metric} \\
&\text{IoT_Infrastructure} \sqsubseteq \exists\ \text{isComposedBy Device} \\
&\text{IoT_Infrastructure} \sqsubseteq \exists\ \text{isComposedBy Link}
\end{aligned}$$

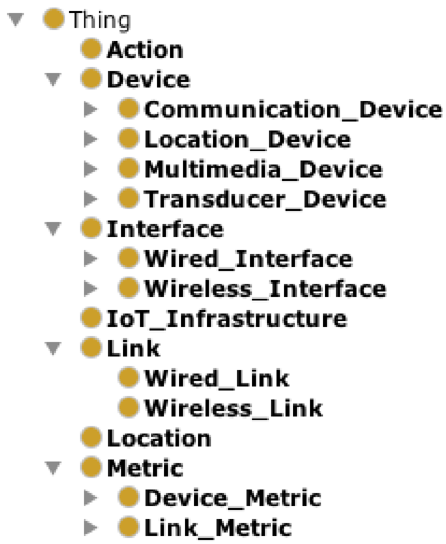

Fig. 1: Class Hierarchy

The composition of the IoT infrastructure is also depicted in Figure 2. Figure 2-a shows how the IoT infrastructure is composed of devices and the communication interfaces and links that they use to exchange information. Additionally, it is of particular interest to store some metrics related to the devices (e.g. Continuity, Downtime, and Packet Loss) and links (e.g. Jitter, Latency) that can be useful for the different management and monitoring mechanisms that are going to be implemented.

On the other side, on Figure 2-b is presented the reversed relationship, or *Object Property*, that indicates that the devices, interfaces, links, and metrics *belongs to* the IoT infrastructure.

For the SusCity project, different devices are considered to take part of the IoT infrastructure, and are grouped in the *Device* class, that is described below. It is important to mention that the classes Device, Metric, Interface, and Link are pairwise disjoint.

$\text{Device} \sqsubseteq \text{Thing}$
$\text{Device} \sqsubseteq \neg\ \text{Metric}$
$\text{Device} \sqsubseteq \neg\ \text{Interface}$
$\text{Device} \sqsubseteq \neg\ \text{Link}$
$\text{Device} \sqsubseteq\ =\ \text{belongsTo IoT\_Infrastructure}$
$\text{Device} \sqsubseteq \exists\ \text{hasDeviceMetric Device\_Metric}$
$\text{Device} \sqsubseteq\ =\ \text{isBackupOf Device}$
$\text{Device} \sqsubseteq \exists\ \text{isPrimaryOf Device}$
$\text{Device} \sqsubseteq \exists\ \text{hasInterface Interface}$

For the *Device* entity, it was particularly important to take into consideration the use of a primary and backup device, that would enable the use of self-healing mechanisms. This was modeled as depicted on Figure 3, with the functional object property *isBackupOf* and its inverse functional object property *isPrimaryOf*, both characterizing that some devices are primaries and some are backups for the aforementioned primary devices.

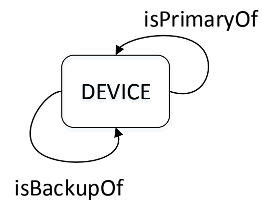

Fig. 3: Device entity model

This will allow instantiating a different backup device in case of a fault in the primary one, whether accidental of provoked. Additionally, the *isBackupOf* property is irreflexive, to avoid that a *Device* becomes backup of itself, rendering to a wrongful solution. The same approach was used for the *Link* class, using the concept of primary and backup links.

The collection of data gathered by the sensors could be automatized or through the use of *Data Loggers*. In the case of the automatized data collection, it is important to include in the model how the devices communicate. The communication interfaces could be wired or wireless. The description of the *Interface* class is provided here.

$\text{Interface} \sqsubseteq \text{Thing}$
$\text{Interface} \sqsubseteq\ =\ \text{belongsTo IoT\_Infrastructure}$
$\text{Interface} \sqsubseteq\ =\ \text{isAttachedTo Device}$
$\text{Interface} \sqsubseteq\ =\ \text{uses Link}$

Another important issue is keeping track of the actual physical location of the devices. This was modeled with the entity *Location* that has the Data Properties: *Latitude*, *Longitude*, and *Elevation*. This relationship, depicted in Figure 4, allows to identify the physical location of the device which will ease the recovery process in case of a failure. A location can host some devices (*hostsDevice*), and a device is located at exactly one physical location (*hasLocation*). Additionally, some Data Properties were also used to model the IoT infrastructure. Table I shows a list summarizing some of the Data Properties and their description is provided.

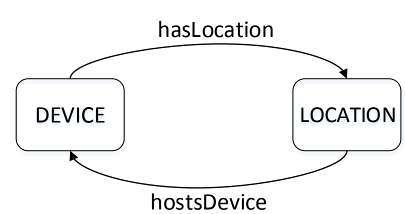

Fig. 4: Device - Location relationship

The final ontology, developed using Protégé [17], was consulted with different members from the SusCity project to verify that it included all the required elements to model the IoT infrastructure to properly use it in the various mechanisms that are going to manage the network. Further validation was carried out and is discussed in the next section.

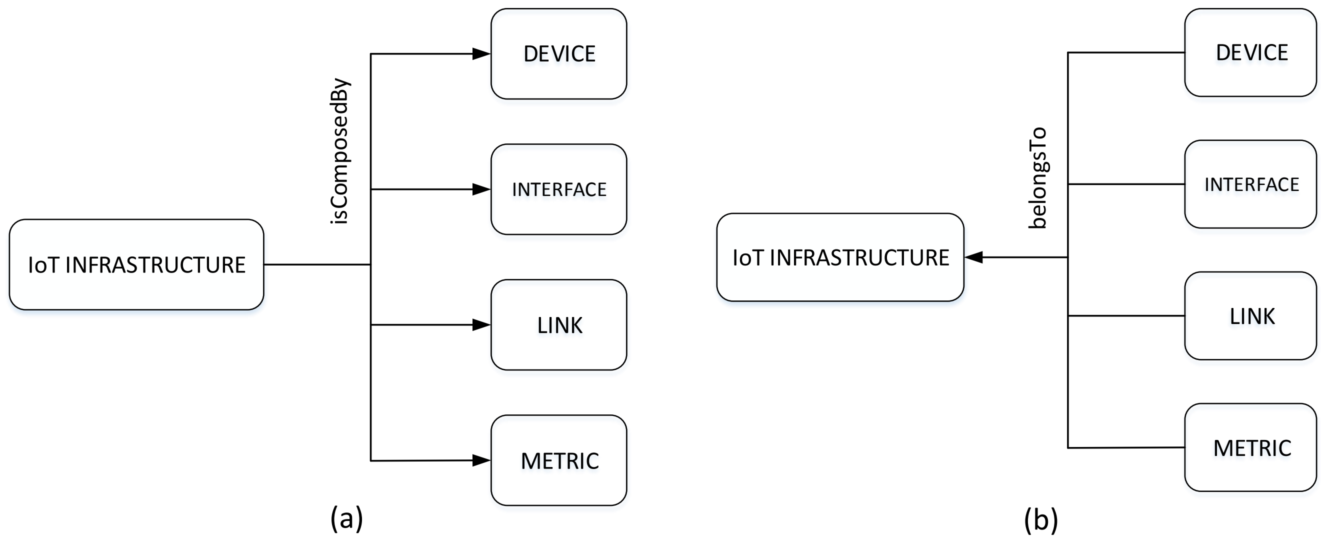

Fig. 2: IoT Infrastructure entity and its Relationships

TABLE I: Data Properties for the SusCity IoT Ontology

Entity	**Data Property**	**Description**
Device	hasIdentifier hasFirmwareName hasFirmwareVersion hasRole hasStatus hasMode hasMeasurementFrequency	Integer representing an unique identifier String describing the name of the firmware String describing the version of the firmware String indicating if the device is primary or backup String indicating if the device status, e.g. up, down String describing the device mode, e.g. power safe Float indicating the time between each capture interval
Link	hasCapacity hasPhyTechnology	Float indicating the link capacity in Mbps String describing the link technology, e.g. fibre, satellite
Metric	hasValue	Float containing the value of the metric measured
Location	hasElevation hasLatitude hasLongitude	Float indicating the elevation in meters of the location Float indicating the latitude of the location Float indicating the longitude of the location
Interface	hasNumberOfAntennas hasPhyAddress	Integer indicating the amount of antennas for a wireless interface HexBinary depicting the physical address of the interface

IV. Evaluating the Ontology

To corroborate the correctness of the ontology, a validation against consistency issues was performed using HermiT 1.3.8.3 [18]. The results obtained showed that the ontology has no consistency issues. Furthermore, DL expressivity metrics were also calculated using Protégé [17], obtaining SHIQ(D). The expressivity evaluation allows obtaining an upper bound on the performance of querying the ontology once it is fully populated. Some restrictions modeled (e.g. inverse object properties) will enable faster response times while querying. The ontology was also evaluated by performing a set of queries. The results are discussed in this section.

A. Querying the Ontology

The ontology was populated with the description of the devices being used in the SusCity project (e.g. smartmeters, sensors) and the communication links that connect them. With this information, some sample queries were designed to verify if the results match the expected values, thus validating the ontology. Since the goal is to use this ontology with different network, service, and application managing mechanisms that are currently under design and development, some possible scenarios were used focused on the resilience of the infrastructure and the latency of the communication network, which are two of the key features to optimize in the network. The scenario used in the tests is described in Figure 5.

The scenario is divided into four zones (Zone 0 to Zone 3) interconnected. Zone 0 has one smartmeter (SM0) with redundant links (L0, L1) connected to two different routers (R0, R1). Zone 1 has two smartmeters (SM5, SM6) connected to R0 and R1 via L4 and L5 respectively. Similarly, Zone 2

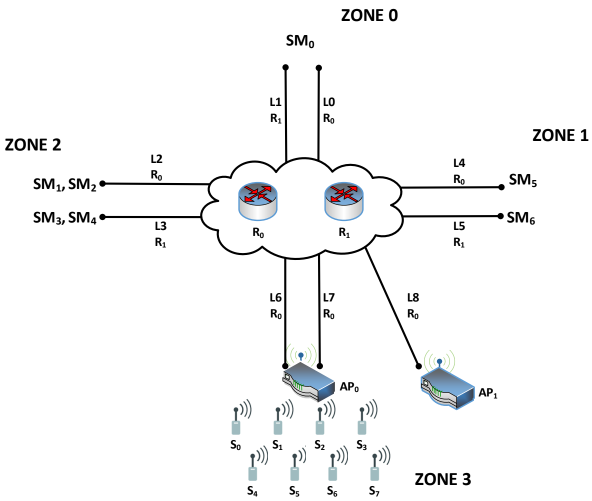

Fig. 5: Testbed used for Validation

has four smartmeters (SM1 to SM4) connected through L2 and L3; and finally, Zone 3 has two access points (AP0, AP1), where AP0 has redundant links to R0 (L6) and R1 (L7) and also is connected to eight CO2 sensors (S0 to S7), and AP1 is connected via L8.

To improve resilience, one common approach is using backup devices that are instantiated once the primary device is down in order to increase the availability. In the test scenario described in Figure 5 some devices were configured as primary while some others were identified as a backup. Additionally, information about the status of the devices (e.g. Up, Down, Rebooting) is also stored.

The first test consisted on querying the ontology to find out which devices are down. The following query provided this information, and the results were as expected according to the information used to populate the ontology. This example allows affirming that this type of queries can be used to build intelligent IoT infrastructure management services, in this case, aimed at the improvement of the resilience.

```
SELECT ?device
WHERE {
?device hasStatus ?status .
FILTER (?status="DOWN")
}
```

This query effectively listed all the devices with status "*Down*". This information is extremely helpful for a mechanism aimed at improving resilience since it clearly identifies the devices that need to be replaced. As stated before, a common technique is to upgrade a backup device to a primary status, thus recovering the failure. To test this, a second query was designed, with the objective of listing the backup devices associated with the primary devices with a current status "*Down*". The status of the backup device was also provided by the query, which is shown below.

```
SELECT ?devbac
WHERE {
?devbac isBackupOfDevice ?devpri .
?devpri hasStatus ?statusp .
?devbac hasStatus ?statusb .
FILTER (?status="DOWN")
}
```

For this query, *devpri* refers to the primary device and *devbac* to the backup device. Similarly, *statusp* and *statusb* refer to the status of the primary and backup devices respectively. Once again, this type of queries could be used on smart management services for the ICT in the IoT.

To further evaluate the ontology other queries were designed, this time with a different management objective in mind. For the following queries, a mechanism to improve the network latency was used as a possible example to test the ontology.

In the scenario designed, the communication links are modeled as well as some metrics of interest such as the jitter. It could be useful to compare the jitter from different links and use this information for instance in a routing mechanism, giving higher priority to links with lower jitter to select these links for the routes. The next query lists the links ranked by the jitter in an ascendant order.

```
SELECT ?link ?metric ?jitter
WHERE {
?link hasLinkMetric ?metric .
?metric sus:hasValue ?jitter .
ORDER BY (?jitter)
}
```

Using the same concept applied to the devices, the links were also modeled as primary and backup. The following query was used to list the backup links whose jitter was smaller than the jitter of its primary correspondent link. This information could be used to update routes and eventually improve the service for final users.

```
SELECT ?linkp ?jitterp ?linkb ?jitterb
WHERE {
?linkb isBackupOfLink ?linkp .
?linkb hasLinkMetric ?metricb .
?linkp hasLinkMetric ?metricp .
?metricb hasValue ?jitterb .
?metricp hasValue ?jitterp .
FILTER (?jitterb < ?jitterp )
}
```

For this query, *linkp* and *linkb* refer to the primary and backup links respectively, *metricp* and *metricb* indicate the metric associated to the primary and backup links, and analogously *jitterp* and *jitterb* specify the jitter for their corresponding links.

Results for all the queries described in this section returned the expected information according to the data that was used to populate the ontology. These results confirm the possibility to use this ontology as input for intelligent management web services, e.g. with a workflow of interaction with the user, where user's queries are translated into SPARQL queries on demand.

V. Conclusions and Future Work

This paper presents an ontology for a Smart City scenario, tailored to the needs of the SusCity project. The ontology comprises as main classes the IoT infrastructure, its devices, communication interfaces and links, and the performance metrics related to them. This information will ease the managing of the infrastructure in order to guarantee its proper work for user services and applications. The ontology was successfully validated with different procedures, including consistency evaluation and also queries to corroborate its correctness.

It is important to mention that although the proposed ontology was designed for the requirements of the SusCity project, it can easily be adapted to other Smart City scenarios since it includes all the main classes involved in a Smart City infrastructure (e.g. sensors, actuators, multimedia device, communication device). For each implementation, it would only be necessary to adjust the population of the ontology in order to comply the different scenarios. This means that the ontology presented in this research could be used in any Smart City scenario just adding the particular description of the given devices related to the specific environment.

For future work, an entire population of the ontology is required. Moreover, it is planned to use this ontology to share information about the status of the infrastructure and its components among the different managing mechanisms for the infrastructure. Finally, it is important to have an automatic mechanism to keep the ontology up-to-date. This will require, for instance, online acquisition of the data describing new devices that are to be dynamically added to the IoT Infrastructure, transforming those descriptions into the ontology's format and enriching the ontology with the new information.

Acknowledgements

This work is financed by national funds via FCT - Foundation for Science and Technology within the scope of projects ≪MITP TB / CS / 0026/2013 - SusCity≫. This work was also co-financed by the DenseNet project ≪PTDC/EEISCR/6453/2014≫, which is a FCT/FEDER/-COMPETE 2020 project.

Note

This is a preliminary version of the paper. The final publication is available at IEEE Xplore via http://ieeexplore.ieee.org/document/7899427

References

[1] IBM, "IBM Smarter Cities - Creating opportunities through leadership and innovation," IBM - https://goo.gl/7aLzQM, 2013. [Online]. Available: https://goo.gl/7aLzQM

[2] R. Petrolo, V. Loscrí, and N. Mitton, "Towards a smart city based on cloud of things," in *Proceedings of the 2014 ACM International Workshop on Wireless and Mobile Technologies for Smart Cities*, August 2014, pp. 61–66.

[3] K. Ashton, "That "Internet of Things" Thing," *RFID Journal*, vol. 1, no. 1, 2009.

[4] G. Lee, P. Jung-Soo, N. Kong, N. Crespi, and I. Chong, "The Internet of Things - Concept and Problem Statement," Working Draft, IETF, Internet-Draft draft-lee-iot-problem-statement-05, July 2012.

[5] SusCity, "SusCity Project," FCT - Foundation for Science and Technology, 2015. [Online]. Available: http://groups.ist.utl.pt/suscity-project

[6] RDFS, "RDF Schema," W3C, 2014. [Online]. Available: https://www.w3.org/TR/rdf-schema

[7] OWL, "OWL 2 Web Ontology Language," W3C, 2012. [Online]. Available: https://www.w3.org/TR/owl2-overview

[8] Turtle, "RDF Turtle," W3C, 2014. [Online]. Available: https://www.w3.org/TR/turtle

[9] SPARQL, "SPARQL," W3C, 2013. [Online]. Available: https://www.w3.org/TR/sparql11-overview

[10] S. De, P. Barnaghi, M. Bauer, and S. Meissner, "Service Modelling for the Internet of Things," in *Proceedings of the Federated Conference on Computer Science and Information Systems (FedCSIS)*, September 2011, pp. 949–955.

[11] J. Byun, K. Seong-Hoon, and K. Daeyoung, "Lilliput: Ontology-Based Platform for IoT Social Networks," in *2014 IEEE International Conference on Services Computing (SCC)*, June 2014, pp. 139–146.

[12] B. Manate, V. Munteanu, and T. Fortis, "Towards a Smarter Internet of Things: Semantic Visions," in *2014 Eighth International Conference on Complex, Intelligent and Software Intensive Systems (CISIS)*, July 2014, pp. 582–587.

[13] S. Hachem, T. Teixeira, and V. Issarny, "Ontologies for the Internet of Things," in *Proceedings of the 8th Middleware Doctoral Symposium*, January 2011, pp. 1–6.

[14] W. Wang, S. De, R. Toenjes, E. Reetz, and K. Moessner, "A Comprehensive Ontology for Knowledge Representation in the Internet of Things," in *IEEE 11th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, 2012.

[15] W. Wang, K. Lee, and D. Murray, "Building a generic architecture for the Internet of Things," in *2013 IEEE Eighth International Conference on Intelligent Sensors, Sensor Networks and Information Processing*, April 2013, pp. 333–338.

[16] D. Perez Abreu and K. Velasquez, "SusCity Ontology," University of Coimbra, 2016. [Online]. Available: https://eden.dei.uc.pt/~dabreu/ontologies/iot_infrastructure.ttl

[17] Protégé, "Protégé," Stanford University, 2015. [Online]. Available: http://protege.stanford.edu

[18] HermiT, "HermiT OWL Reasoner," University of OXFORD - Information Systems Group, 2015. [Online]. Available: http://www.hermit-reasoner.com