



HAL
open science

IoT-O, a Core-Domain IoT Ontology to Represent Connected Devices Networks

Nicolas Seydoux, Khalil Drira, Nathalie Jane Hernandez, Thierry Monteil

► **To cite this version:**

Nicolas Seydoux, Khalil Drira, Nathalie Jane Hernandez, Thierry Monteil. IoT-O, a Core-Domain IoT Ontology to Represent Connected Devices Networks. Knowledge Engineering and Knowledge Management: 20th International Conference, EKAW 2016, Nov 2016, Bologne, Italy. pp.561 - 576, 10.1007/978-3-319-49004-5_36 . hal-01467853v1

HAL Id: hal-01467853

<https://hal.science/hal-01467853v1>

Submitted on 14 Feb 2017 (v1), last revised 25 Apr 2019 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

IoT-O, a core-domain IoT ontology to represent connected devices networks

Nicolas Seydoux^{1,2,3}, Khalil Drira^{2,3}, Nathalie Hernandez¹, Thierry Monteil^{2,3}

¹ IRIT Maison de la Recherche, Univ. Toulouse Jean Jaurès,
5 allées Antonio Machado, F-31000 Toulouse
{nseydoux,hernande}@irit.fr

² CNRS, LAAS, 7 avenue du Colonel Roche,
F-31400 Toulouse, France
{nseydoux,khalil,monteil}@laas.fr

³ Univ de Toulouse, INSA, LAAS, F-31400, Toulouse, France

Abstract. Smart objects are now present in our everyday lives, and the Internet of Things is expanding both in number of devices and in volume of produced data. These devices are deployed in dynamic ecosystems, with spatial mobility constraints, intermittent network availability depending on many parameters (e.g. battery level or duty cycle), etc. To capture knowledge describing such evolving systems, open, shared and dynamic knowledge representations are required. These representations should also have the ability to adapt over time to the changing state of the world. That is why we propose IoT-O, a core-domain modular IoT ontology proposing a vocabulary to describe connected devices and their relation with their environment. First, existing IoT ontologies are described and compared to requirements an IoT ontology should be compliant with. Then, after a detailed description of its modules, IoT-O is instantiated in a home automation use case to illustrate how it supports the description of evolving systems.

1 Semantic interoperability, a challenge for the IoT

The Internet of Things (IoT) is gaining more and more traction: some projectionists predict up to 50 billion devices connected in the next five to ten years [5]. The Things of the IoT allow to connect the physical world and virtual representations. IoT applications are based on very heterogeneous devices and technologies, and are deployed in domains as diverse as agriculture, domotics⁴, smart cities or e-health. Two types of interoperability issues can be identified: syntactic and semantic, brought by the variety of domains and data models[6]. This paper focuses on semantic interoperability, the ability of systems to attribute the same meaning to the data they exchange.

Semantic interoperability is based on shared, unambiguous, machine-understandable vocabularies, which is why semantic web principles and technologies

⁴ Home automation

are seen as semantic interoperability providers, as [14] expresses for the specific domain of IoT. Knowledge expressed in open formats can be shared and reused, and ontologies can evolve to adapt to new contexts or usages. To ensure the reusability of semantic models across projects and domains, good practices in ontology design have been proposed. In IoT projects, many ontologies have been built, but not always according to these guidelines, hence limiting their reusability. This is why we propose IoT-O⁵, an IoT core-domain modular ontology engineered for reusability and extensibility. IoT-O is also available on the LOV⁶.

IoT systems are strongly bound to their environment, because they are composed of devices in contact with the physical world. Sensors collect data about their environment, and actuators are devices performing actions that have a direct impact on the world: light bulbs, motors, air conditioning, etc. This paper aims at showing how IoT-O can be used as an ontology to semantically describe devices and data in order to make systems aware of their environment, its evolution, and the changes they can bring to it. Such a description allows smart agents to transform their environment thanks to connected actuators, according to the perceptions they have of it through connected sensors.

In the remainder of this paper, section 2 introduces a motivating use case that will serve to instantiate portions of IoT-O. Section 3 presents the design process of IoT-O, and gives an overview of the ontology. Finally, section 4 details how IoT-O is instantiated in the use case.

2 Motivating use case

IoT technologies can have a direct impact on the everyday life of citizens, since it connects their physical environment to virtual applications. That is especially relevant in the case of domotics, where the home can be equipped with multiple low-power devices to provide new services. At LAAS-CNRS, the ADREAM project⁷ aims at conducting research on smart buildings thanks to an instrumented, energy-positive building. It is equipped with more than 4500 sensing devices, producing up to 500,000 measures a day. Inside the building, there is a mock-up apartment equipped with commercial devices from various vendors. Deployed devices include sensors (temperature, luminosity, humidity, pressure), actuators (fan, space heater, multiple lamps), which communicate using different technologies (phidget, ethernet, zigbee) with gateways connected to a central server.

However, small highly distributed devices usually have a limited processing power, which restrict the range of applications they can support. More complex agents can interact with these devices to collect their data and perform advanced processing to provide a higher level service. In our use case, centered on an elderly healthcare scenario, the complex agent is a robot. It is present in the house, and

⁵ <http://www.irit.fr/recherches/MELODI/ontologies/IoT-O>

⁶ <http://lov.okfn.org/dataset/lov/vocabs/ioto>

⁷ <http://www.laas.fr/public/en/adream>

performs tasks such as helping the person in case of fall, moving heavy objects, pushing a wheelchair, fetching objects and bringing medications. Some of these tasks require the robot to know **where the person is** in the apartment. To have this information, the robot can move around the apartment, scan it with its embedded cameras, and through image processing figure out where the person is. However, it requires the robot either to follow the person around all the time, or to scan the apartment completely each time it has to find the person. To make the robot more acceptable to the person, the house can be equipped with an IoT system, collecting information useful to the robot, such as information given by **presence sensors**. Moreover, the connected devices can provide new functionalities to the robot: he can easily interact with connected light switches or sensors. Our use case is composed of two scenarios: the robot must bring pills at fixed hours to the person using the presence sensors to locate her, and the robot must control the temperature in the apartment using temperature sensors and connected fans to improve the comfort of the person.




Fig. 1. PR2, the companion robot

In this use case, both syntactic and semantic interoperability are required, among the devices and between the devices and the robot. Syntactic interoperability is ensured using OM2M⁸, an open-source horizontal integration platform implementing the oneM2M⁹ standard. On top of OM2M, another platform, SemIoTics, enriches the collected data with semantic descriptions, and makes them available to the robot through a REST interface. SemIoTics is driven by a knowledge base capturing knowledge about the devices of the system represented according to our core-domain IoT ontology, and about the environment shared by the robot and the devices (here, the apartment). It is a Java software developed to showcase the role of semantic web technologies in IoT data management, based on Apache Jena. The robot itself is also a semantically enabled

⁸ om2m.org

⁹ <http://onem2m.org/>

agent, it uses a "common sense" ontology and a knowledge base to reason about its 3D environment, as described in [13]. The knowledge specific to the robot relies on ontologies out of the scope of this paper, but its knowledge base can be extended with any ontology, including IoT-O (as it is done in this paper).

That is why the knowledge described in this paper is implemented in a dedicated knowledge base using IoT-O, ADREAM-Robot¹⁰: the ontology is shared by the robot and SemioTics, and each system has its own knowledge base. The synchronization between the knowledge bases of the different agents is out of the scope of this paper. The use case focuses on home automation, but IoT-O and our approach are generic enough to be adapted to other domains. For instance, it could be used to support an air quality monitoring system in a smart city, by describing the sensors that collect the data and the services the citizens can subscribe to. The usage of IoT-O and its module in the use case is double: it is used to model the observations about modifications of the apartment, allowing the robot to keep an up-to-date representation of its environment, but also to model the changes the robot wants to make into the apartment through its actions and through the connected devices.

3 IoT-O, not just another IoT ontology

The design of IoT-O is compliant with the NeOn methodology, presented in [3].

The first step of the NeOn process is to define requirements. We split them in two types: **conceptual**, regarding the concepts that should be present in the ontology (detailed in section 3.1), and **functional**, regarding the ontology structure and design principles (detailed in section 3.2).

These requirements are used to analyze existing IoT ontologies: Semantic Sensor Network (SSN)¹¹, Smart Appliance REference (SAREF)¹², iot-ontology¹³, IoT-lite¹⁴, Spitfire¹⁵, IoT-S¹⁶, SA¹⁷ and the oneM2M base ontology¹⁸. These ontologies are IoT ontologies for which we have found information on the web. Further details are available on the Linked Open Vocabularies for the IoT (LOV4IoT)¹⁹, a recent initiative that lists IoT ontologies, even if they are not referenced on the LOV because they fail to comply with its requirements recalled in [6]. Ontologies related to specific domains impacted by IoT (domotics, agriculture, smart cities...) are out of the scope of this study.

As recommended by NeOn, reusable ontologies that are compliant with parts of the requirements are integrated in our design process. They are analyzed and

¹⁰ <https://www.irit.fr/recherches/MELODI/ontologies/Adream-Robot>

¹¹ <http://purl.oclc.org/NET/ssnx/ssn>

¹² <http://sites.google.com/site/smartappliancesproject/ontologies>

¹³ <http://ai-group.ds.unipi.gr/kotis/ontologies/IoT-ontology>

¹⁴ <http://iot.ee.surrey.ac.uk/fiware/ontologies/iot-lite>

¹⁵ <http://sensormeasurement.appspot.com/ont/sensor/spitfire.owl>

¹⁶ <http://personal.ee.surrey.ac.uk/Personal/P.Barnaghi/ontology/OWL-IoT-S.owl>

¹⁷ http://sensormeasurement.appspot.com/ont/sensor/hachem_onto.owl

¹⁸ http://www.onem2m.org/ontology/Base_Ontology/

¹⁹ <http://www.sensormeasurement.appspot.com/?p=ontologies>

presented in section 3.3. The core-domain ontology we propose is then described in section 3.4.

3.1 The core concepts of IoT

Conceptual requirements These requirements come from an analysis of the IoT domain, driven by the home automation use case introduced in section 2, but not limited to it: the use case is not seen as an end per se, but as an instantiation of the general domain of the IoT. To be reusable in a wide scope of domains, an IoT ontology should contain a set of key concepts. These are representative of IoT systems with no regard to the application domain. This approach facilitates the merging of data collected in different domains for horizontal applications, and allows the ontology to be an extendable core-domain ontology. We distinguish namely:

- **”Device”** and **”software agent”** constitute the two basic components of an IoT system, composed of both physical and virtual elements. The devices can be of two principle types, not mutually exclusive, that are listed below.
- **”Sensor”** are devices acquiring data, and **”observation”** describe the acquisition context and the data collected by the system. These concepts capture the perception the system has of the evolutions of its environment.
- **”Actuator”** are the devices that enable the system to act on the physical world, and **”action”** represents what they can perform. These concepts capture the knowledge the system has on its own abilities to impact its environment, and to make it evolve.
- **”Service”**: In many cases, the IoT and the programmable web are very close. Connected devices can be seen as service providers and consumers, and by specifying a notion of service, every aspect of an IoT system can be represented.
- **”Energy”**: In the paradigm of pervasive computing, many distributed Things perform computations. Most of these Things being physical devices, a complete modelling of the system will include a description of their energy consumption. Energy management is a crucial topic in IoT systems.
- **”Lifecycle”**: Be it data, devices or services, IoT components are all included in different scales of lifecycles. Devices are switched on and off, services are deployed or updated, pieces of data become outdated... The evolution through a set of discrete states representing a lifecycle is an important concept for IoT systems.

Concept coverage by existing ontologies Table 1 sums up the assessment of existing IoT ontologies regarding the presence of key concepts. One star means that the concept is superficially represented (coarse-grained specialization, few data/object properties), two stars that the requirement is covered, and stars between parentheses indicate that the requirement is met by an included ontology. IoT-O, the ontology we propose, is also included for comparison. Note that we focus on connected device ontologies, and exclude, on purpose, the ontologies SSN is based upon, since they are only focused on sensors and observation,

which is only a subset of the identified key concepts. We can observe that some of the IoT ontologies cover most of the key concepts but none of them covers them all. Moreover, the different concepts are not represented with the same level of expressivity. In *iot-ontology* and SAREF, key concepts such as Actuator or Action are present but their representation is limited. For example, an actuator is defined as a device that modifies a property. This is less expressive than what can be expressed for a sensor with SSN which proposes a deep modeling of the sensors and the property they observe, but also of the relations between the sensors and their observations, and of the observations themselves. In eDI-ANA²⁰, an ontology referenced by SAREF, some specializations of actuator are given, but the mappings from these specializations to the *saref:Actuator* concept are not available directly. This analysis highlights the fact that an ontology for Actuators and Actions is needed (c.f. section 3.3). This analysis also highlights the failure of existing IoT ontologies in representing correctly all IoT key concepts. As these concepts are not limited to the IoT domain, reusing ontologies dedicated to them (such as SSN for sensor) could help gain in expressivity, as is shown in section 3.2.

Table 1. Key concept coverage in IoT ontologies

	Structured by ODP	Modular	Reuses external ontologies	Aligned with upper ontologies	One the LOV	Available online
<i>iot-ontology</i>			*	**	N	Y
<i>saref</i>		**	*		Y	Y
OWL-IoT-S	(*)	*	**	*	N	Y
SA	(*)	*	**	**	N	N
<i>iot-lite</i>					N	Y
<i>spitfire</i>			*	**	Y	N
<i>ssn</i>	**	**	*	**	Y	Y
<i>oneM2M</i>					N	Y
IoT-O	(**)	**	**	**	Y	Y

3.2 Good practices for ontology design

Functional requirements These requirements capture ontology design guidelines and general semantic web good practices in a domain-agnostic fashion.

Reusability: One of the most important aspects of an ontology in such a broad domain as IoT is reusability: if an ontology is ad-hoc to a project, the work done in its definition will not benefit further projects. It is a critical issue that can be solved by different, non-mutually exclusive approaches:

- **Modularization:** as stated in [1], designing ontologies in separated modules makes them easier to reuse and/or extend. IoT applications are related to many various domains, and it is difficult to capture all these application domains in the same ontology. Modular ontologies can be combined together according to specific needs, which is a more scalable approach.

²⁰ <https://sites.google.com/site/smartappliancesproject/ontologies/ediana-ontology>

- **Ontology Design Patterns:** were introduced in [4]. Designing ontologies that respect Ontology Design Pattern (ODP) increases reusability and their potential for alignment, as shown in [15]. ODPs capture modelling efforts: using them is a way to capitalize on previous work, and to take advantage of the maturity of the semantic web compared to the IoT.
- **Reuse of existing sources:** avoids redefinition, and prevents from having to align a posteriori the redefined concepts to the existing sources for interoperability. It is a key requirement for interoperability, which is a real issue in heterogeneous systems.
- **Alignment to upper ontologies:** Upper-level ontologies define very abstract concepts in a horizontal manner. They articulate very diverse domain-specific ontologies, which is crucial for broad domains like IoT.
- **Compliance with the LOV requirements:** The LOV²¹ is an online vocabulary register that increases visibility of vocabularies, and favours reuse by ensuring the respect of good practices listed in [6].

Level of formalism: To use the full advantages of the semantic description of devices and data, the description should enable reasoning and inference. This choice is motivated by the possibilities it opens:

- Applied to data, it is a way to bring context-awareness, as presented in [8]
- Applied to devices, it enables Thing discovery or self-configuration [2]
- Applied to services it enables automatic composition as in [7]

However, for concrete applications, the model should also be decidable, and in reasonable time, which de facto excludes an OWL-full model: OWL-DL is therefore the best choice. All surveyed ontologies are expressed in OWL-DL.

Table 2. Reusability of IoT ontologies

	Structured by ODP	Modular	Reuses external ontologies	Aligned with upper ontologies	One the LOV	Available online
iot-ontology			*	**	N	Y
saref		**	*		Y	Y
OWL-IoT-S	(*)	*	**	*	N	Y
SA	(*)	*	**	**	N	N
iot-lite					N	Y
spitfire			*	**	Y	N
ssn	**	**	*	**	Y	Y
oneM2M					N	Y
IoT-O	(**)	**	**	**	Y	Y

Assessment of existing IoT ontologies Table 2 shows that the semantic web best practices for reusability are not always followed: some ontologies are not available online, and the majority is not compliant with the requirements of the LOV. External ontologies are generally not reused, with the exception of SSN. OWL-S, a service ontology is reused in only one case. The other surveyed

²¹ <http://lov.okfn.org>

ontologies propose redefinitions of the service concept. For example, SAREF redefines the concepts present in multiple ontologies, and proposes alignments in an external, textual document. Design patterns have only been used in ontologies importing SSN. Upper ontologies used are DUL²² (especially used by SSN) and SWEET²³ (for SA). The limited reuse of ontologies shows a lack of federating ontologies, apart from SSN. SSN being a modular ontology compliant with the semantic web good practices, it is possible to say that these guidelines favour reuse. Section 3.3 focuses on such good practices.

3.3 Reused ontologies for IoT-O

Identification of existing ontologies is included in the NeOn process. Some concepts, which are part of the conceptual requirements are defined by existing ontologies that are imported in IoT-O to avoid redefinition. SSN is a widely used W3C recommended ontology for sensors and observations. To define the notion of service, IoT-O imports Minimal Service Model (MSM), a lightweight service ontology which is generic enough to represent both REST and WSDL services (contrary to OWL-S²⁴). The notion of energy consumption dedicated to the IoT is specified in PowerOnt, an ontology referenced by SAREF. The concepts of lifecycle are described using Lifecycle²⁵, a lightweight vocabulary defining state machines. We extended Lifecycle in the IoT-lifecycle²⁶ ontology with classes and properties specific to the IoT. Finally, to maximize extensibility and reusability, IoT-O imports DUL²⁷, a top-level ontology, and aligns all its concepts and imported modules with it.

Focus on SAN: However, no ontology describes the concept of actuator the way SSN describes the concept of sensor. This is why we propose the Semantic Actuator Network (SAN)²⁸ ontology. Actuators are devices that transform an input signal into a physical output, making them the exact opposite of sensors. SAN is built around Action-Actuator-Effect (AAE)²⁹, a design pattern we propose, inspired from the Stimulus Sensor Observation (SSO) design pattern described in [9]. Fig. 2 shows a representation of both the AAE and the SSO design patterns. SSN models the state of the world through stimuli converted by sensors into abstract observations, making the system able to be aware of the evolution of its environment. SAN is complementary: it models the transformation of abstract actuations by actuators into real-world effects, leading to the representation of the evolution the system brings into its environment. Further details of the main classes of SAN are provided in section 4.4.

²² <http://www.ontologydesignpatterns.org/ont/dul/DUL.owl>

²³ <http://sweet.jpl.nasa.gov/>

²⁴ <https://www.w3.org/Submission/OWL-S/>, more dedicated to WSDL-based services

²⁵ <http://vocab.org/lifecycle/schema>

²⁶ <https://www.irit.fr/recherches/MELODI/ontologies/IoT-Lifecycle>

²⁷ <http://www.ontologydesignpatterns.org/ont/dul/DUL.owl>

²⁸ <https://www.irit.fr/recherches/MELODI/ontologies/SAN>

²⁹ <http://ontologydesignpatterns.org/wiki/Submissions:Actuation-Actuator-Effect>

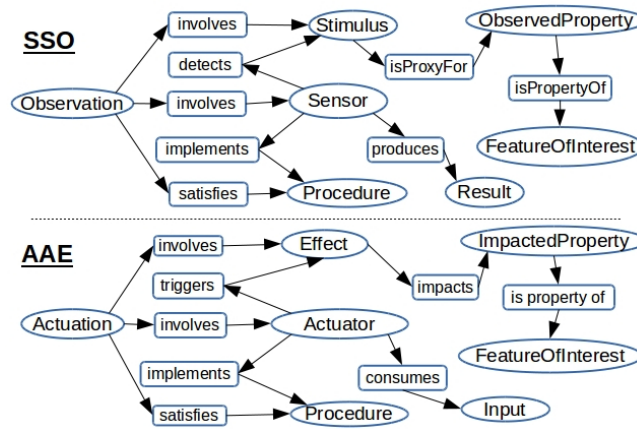


Fig. 2. The SSO and the AAE design patterns, structuring respectively SSN and SAN

3.4 IoT-O, a modular core-domain IoT ontology

IoT-O, the core-ontology we propose is composed of several modules. IoT-O’s architecture is summarized in figure 3. The names of the newly created resources are in red and highlighted, the names of the reengineered resources are underlined, and the arrows show dependencies. Solid arrows represent imports, and dashed arrows the reuse of concepts without import.

The modules of IoT-O:

- The **Sensing module** describes the input data. Its main classes come from SSN: *ssn:Sensor* and *ssn:Observation*. *ssn:Device* and its characteristics (*ssn:OperatingRange*, *ssn:Deployment...*) provide a generic device description.
- The **Acting module** describes how the system can interact with the physical world. Its main classes come from SAN: *san:Actuator* and *san:Actuation*. It also reuses SSN classes that are not specific to sensing, such as *ssn:Device*.
- The **Lifecycle module** models state machines to specify system life cycles and device usage. Its main classes are *lifecycle:State* and *lifecycle:Transition*.
- The **Service module** represents web service interfaces. Its main classes come from MSM: *msm:Service* and *msm:Operation*. Services produce and consume *msm:Messages*, and RESTful services can be described with hRest.
- **Energy module:** IoT-O’s energy module is defined by PowerOnt. It provides the *poweront:PowerConsumption* class, and a set of properties to express power consumption profiles for appliances.

The core of IoT-O: IoT-O³⁰ is both the name of the ontology and of the top module. It gives a conceptualization of the IoT domain, independent of the application, providing classes and relationships to link the underlying modules. Since

³⁰ <http://www.irit.fr/recherches/MELODI/ontologies/IoT-O.owl>

in four steps : Monitoring, Analysis, Planning and Execution. The K stands for Knowledge, because the behaviour of the autonomic agent at each step of the loop is guided by a knowledge base, in the general meaning of the term (including but not restricted to the W3C's formalisms of knowledge representation).

In this use case, SemIoTics is performing the Monitoring and the Execution steps when connected devices are involved, and the robot performs the Analysis and the Planning, as well as part of the Monitoring and Execution steps. The robot and SemIoTics have distinct knowledge bases, even if in fig. 4, a unique knowledge base is represented as the two systems exchange knowledge freely through a rest interface. Consistence issues are not considered in this work, as only one smart agent interacts with the system.




Fig. 4. A representation of the MAPE-K loop, split between the robot and SemIoTics

The process described in figure 4 structures the use case: data is first gathered by the sensors, and enriched by SemIoTics. The enriched observations are processed by the robot, which decides to perform actions represented as enriched actuations. These actuations are sent to SemIoTics, which translates them into raw commands for the actuators to perform. In complement to IoT-O, the dogont³¹ ontology is used to describe the apartment and the location of devices inside it. Dogont is an ontology identified in the SAREF project, and it is imported by Poweront. We aligned it to IoT-O to integrate it to the use case.

4.2 Monitoring, where raw sensor data become meaningful observations

The first step of the MAPE-K loop is the monitoring of the controlled system. In the apartment, sensors produce data reflecting their observations. This data is enriched to become a reusable piece of knowledge. Enrichment of sensor data is performed using the SSN ontology, which is in the Sensing module of IoT-O.

³¹ <http://elite.polito.it/ontologies/dogont/dogont.html>

Each *ssn:Sensor* has an *ssn:Observation* stream composed of *ssn:SensorOutput* whose value is described by *ssn:ObservationValue*. For provenance purposes, a *ssn:SensorOutput* can be linked to its original representation (before enrichment) with the *iot-o:hasRawRepresentation* data property. The sensor’s characteristics (*ssn:MeasurementProperty*, the *ssn:Property* of the *ssn:FeatureOfInterest* it observes) are used to enrich the observation as well. IoT-O and SSN are generic ontologies, so they might need to be extended with application-specific modules to be fully functional. Such extension is proposed in the Adream-Robot module³². The *ssn:Observation* allow the representation of a characteristic of the environment at a given point in time. The temporality of the sensor measures (and of actuators actions) are represented by a *san:hasDateTime* relations with a *http://w3c.org/2006/time#Instant*, itself characterized by an *iot-o:hasTimestamp* data property. All the observations related to the same point in time are connected to the same individual, allowing the agent to have a timed representation of its environment and of its evolution.

In our use case, presence sensors and a temperature sensors produce raw observations in the form of XML documents standardized according to the oneM2M Content Instance resource type. The enrichment process requires an approach specific to the data, either by writing a dedicated enrichment script, or by using semantic annotations in the data as in [12], where raw data is stored in relational databases and the database schema is annotated for enrichment. SemioTics uses a dedicated enrichment script that could in the future be extended by producing annotated data.

The presence observation indicates the position of the person in the apartment, and the temperature observation measures the temperature at a given point in space and time, both in the form of *ssn:ObservationValue* instances. This enriched information is accessed by the robot through SemIoTics’ REST interface, and it is used to update the robot’s representation of the world. This representation of the world is stored in the robots knowledge base, and used as a context in the Analysis step.

4.3 Analysis: aggregation of observations in abstract symptoms

In the Analysis step, the robot processes his own representation of the world to determine high-level symptoms that need to be addressed by actions.

In the medication scenario, the robot compares the present time to the time when the medication is due to generate the symptom ”Medication must be delivered” if necessary.

In the temperature control scenario, user preferences are represented using the concepts defined in yet another module: Autonomic³³. *ssn:Property* of the environment controlled by the robot within explicit boundaries expressed in the form of *autonomic:PropertyConstraints* are classified as *autonomic:Constrained-Property*. In our use case, the *ssn:Property* temperature of the *ssn:FeatureOfInterest* living room air has two constraints, instances of *autonomic:Maximum-*

³² <https://www.irit.fr/recherches/MELODI/ontologies/Adream-Robot>

³³ <http://www.irit.fr/recherches/MELODI/ontologies/Autonomic>

Value (25°C) and *autonomic:MinimumValue* (19°C). The last *ssn:Observation-Value* of the *autonomic:ConstrainedProperty* is out of the bounds defined by the *autonomic:PropertyConstraint* (26°C instead of 25), so the temperature is classified by the reasoner as an *autonomic:OutOfBoundsProperty* thanks to custom rules.

4.4 Planning, where symptoms are used to create a plan

In the planing phase, the autonomic agent uses the inferred symptoms and policies defined by the user or by the administrator beforehand to define a series of actions that have to be implemented on the system.

In the medication scenario, the robot uses its representation of its environment to locate the person, as it is kept updated in the monitoring phase thanks to the knowledge produced by the sensors and SemioTics. The robot will plan a trajectory to fetch the medication and to reach the person. In this case, the representation of the trajectory itself is ad-hoc to the robot, and isn't linked to IoT-O. The ontology is used to connect the robots internal representation of the world with the observations collected by the sensors and enriched by SemIoTics, providing semantic interoperability between the robot and SemIoTics. If the robot expresses its needs using the same ontology as SemIoTics, or if their ontologies are aligned, it can seamlessly use elements measured by the sensors to plan its trajectory.

In the temperature control scenario, the description of the actions is performed using SAN, the actuator ontology that also describes the actuators in the system. The agent, with successive queries to the knowledge base, will look for *san:Actuator* instances that *san:actsOn* the *autonomic:OutOfBoundsProperty*, and which *autonomic:ImpactOnProperty* is coherent with the symptom. In the example, since the temperature is too high, the *adream-model:fan* can be used, but also the *adream-model:spaceHeater*, since its *adream-model:turnOff* operation has a *adream-model:NegativeImpact* on the temperature. The orchestration of these actions (if need be) are determined using the Lifecycle module of IoT-O, which represents the devices as state machines by integrating the Objects with States (ows)³⁴ ontology design pattern. *ssn:Device* (superclass of both *ssn:SensingDevice* and *san:ActuatingDevice*) are objects that *ows:hasState exactly 1 ows:State*, because objects should only be in one state at a time. The *ows:State* is equivalent to the *lifecycle:State* (from the Lifecycle³⁵ vocabulary, extended by the IoT-Lifecycle³⁶ ontology), and *lifecycle:State* are connected by *lifecycle:Transition* instances. Thanks to this vision of state machines, stateful transitions (that are only available in certain states of the device) can be represented. Only *msm:Operation* instances that *iot-o:isGroundedBy* a *san:-Actuation* that *iot-lifecycle:triggersTransition* a *lifecycle:Transition* that is a *lifecycle:possibleTransition* of the device current *lifecycle:State* can be called at a given time. For instance, the fan *adream-model:turnOff* operation will only be

³⁴ <http://delicias.dia.fi.upm.es/ontologies/ObjectWithStates.owl>

³⁵ <http://vocab.org/lifecycle/schema>

³⁶ <http://www.irit.fr/recherches/MELODI/ontologies/IoT-Lifecycle>

available if the space heater is on. In our example it is off, so the agent plans to turn on the fan and creates the corresponding *san:ActuationValue*. The selection of devices and their operations is driven by necessity (only the devices impacting the right property are selected), but it can also be driven by policies based on knowledge about the devices intrinsic characteristics expressed with *san:ActuatingCapability*, composed of *san:ActuatingProperty* that create an actuator profile. It can be used to minimize energy consumption (combined with the Energy module), to optimize reaction time...

4.5 Execution, where the plan is converted into actions

In the execution step, the robot implements the planned actions.

For the medication scenario, the robot fetches the medication and brings it directly to the person, it doesn't have to search for her in the apartment. The MAPE-K loop can be repeated while the robot is moving to update the trajectory if the person moves in the house.

For the temperature control scenario, the robot transmits the *san:ActuationValue* that it wants the system to implement to SemIoTics via a REST interface. SemIoTics will handle the transformation of the knowledge into a representation that can be processed by the target device. This translation can be driven by the semantic description of *msm:Operations*, or dedicated annotations as in [11], where XML schemas are annotated for transformation from RDF to XML. SemIoTics uses the semantic description of operations to perform lowering, and perspectives for this technique are presented in section 5. This translation enables the interaction with low-level, constrained devices that are not able to process complex knowledge representations.

5 Conclusion and future works

This paper introduces IoT-O, a modular core-domain IoT ontology designed to be compliant with identified requirements. After a detailed presentation of its modules, an instantiation of IoT-O is presented in a home automation use case. IoT-O is used to bring semantic interoperability between SemIoTics, a platform enabling semantic access to connected devices, and a robot. SemIoTics and the robot implements the MAPE-K loop, an autonomic computing pattern, and uses IoT-O at each step of the loop to describe knowledge about the connected devices and about the data they produce and consume. The ontology describes the evolving state of the robot's environment through sensor observations, and the capabilities the system offers to impact this environment through the devices.

In this paper, enrichment and lowering techniques (allowing the transformation back and forth from data to knowledge) have been overviewed. Such techniques are essential to include constrained devices into the IoT: enriched data is more reusable than raw data, but it is heavier to exchange and process, so transformation is required between the end devices (sensors, actuators) and the more powerful nodes of the IoT, e.g. gateways, servers, laptops... We are currently working on such an approach. Other perspectives of our work will be

to manage data flows over time in order to learn from previous decisions and their consequences to produce explicit knowledge and enrich policies, and synchronization of the distributed knowledge bases of the smart agents: compared to the use case, multiple agents should be able to exchange knowledge about their environment and to maintain coherence between their representations of the world.

References

1. Aquin, M.: Modularizing ontologies. *Ontology engineering in a networked world* Springer B, 9–34 (2012)
2. Chatzigiannakis, I., Hasemann, H., Karnstedt, M., Kleine, O., Kröller, A., Leggieri, M., Pfisterer, D., Römer, K., Truong, C.: True Self-Configuration for the IoT. In: 3rd International Conference on the Internet of Things (IOT) (2012)
3. del Carmen Suarez de Figueroa Baonza, M.: NeOn methodology for building ontology networks : specification, sheduling and reuse. Ph.D. thesis (2010)
4. Gangemi, A.: *Ontology Design Patterns for Semantic Web Content*. *History* 3729(4), 262–276 (2005)
5. Ganz, F., Puschmann, D., Barnaghi, P., Carrez, F.: A Practical Evaluation of Information Processing and Abstraction Techniques for the Internet of Things. *IEEE Internet of Things Journal* 2(4), 340–354 (2015)
6. Gyrard, A., Serrano, M., Ateazing, G.A.: Semantic web methodologies, best practices and ontology engineering applied to Internet of Things. In: 2015 IEEE 2nd World Forum on Internet of Things (WF-IoT). pp. 412–417. IEEE (2015)
7. Han, S.N., Lee, G.M., Crespi, N.: Towards Automated Service Composition Using Policy Ontology in Building Automation System. In: 2012 IEEE Ninth International Conference on Services Computing. pp. 685–686 (2012)
8. Henson, C., Sheth, A., Thirunarayan, K.: Semantic perception: Converting sensory observations to abstractions. *IEEE Internet Computing* 16(2), 26–34 (2012)
9. Janowicz, K., Compton, M.: The Stimulus-Sensor-Observation Ontology Design Pattern and its Integration into the Semantic Sensor Network Ontology. In: Proceedings of the 9th International Semantic Web Conference, 3rd International Workshop on Semantic Sensor Networks. pp. 7–11 (2010)
10. Kephart, J., Chess, D.: The vision of autonomic computing. *Computer* 36(1), 41–50 (jan 2003)
11. Kopecký, J., Vitvar, T., Bournez, C., Farrell, J.: SAWSDL: Semantic Annotations for WSDL and XML Schema. *IEEE Internet Computing* 11(6), 60–67 (2007)
12. Le-Phuoc, D., Quoc, H., Parreira, J.X., Hauswirth, M.: The linked sensor middleware—connecting the real world and the semantic web. In: *Semantic Web Challenge 2011*. pp. 1–8. No. April 2005 (2011)
13. Lemaignan, S.: *Grounding the Interaction: Knowledge Management for Interactive Robots*. Ph.D. thesis (2012), <http://d-nb.info/1030099995/34>
14. Murdock, P.: *White paper: Semantic interoperability for the web of things* (2016), <http://dx.doi.org/10.13140/RG.2.2.25758.13122>
15. Scharffe, F., Euzenat, J., Fensel, D.: Towards design patterns for ontology alignment. In: *Proceedings of the 2008 ACM symposium on Applied computing - SAC '08*. p. 2321. ACM Press, New York, New York, USA (mar 2008)