Semantics in the Edge: Sensors and Actuators in the Web of Linked Data and Things

Semantic Web - (2020) 0
IOS Press

Semantics in the Edge: Sensors and Actuators in the Web of Linked Data and Things

Editorial

Federica Cena [a] and Armin Haller [b] and Maxime Lefrançois [c,*]
[a] *Dipartimento di Informatica, Università degli Studi di Torino, Turin, Italy*
E-mail: federica.cena@unito.it
[b] *Research School of Management & Research School of Computer Science, Australian National University, Canberra, Australia*
E-mail: armin.haller@anu.edu.au
[c] *Univ Lyon, MINES Saint-Étienne, CNRS, Laboratoire Hubert Curien UMR 5516, Saint-Étienne, France*
E-mail: maxime.lefrancois@emse.fr

Editors: Pascal Hitzler, Wright State University, USA; Krzysztof Janowicz, University of California, Santa Barbara, USA

Keywords: Internet of Things (IoT), Web of Things (WoT), Cyber-Physical Systems, Ontology, Knowledge Representation, semantics, linked data

1. Introduction

The rapid advancement and ubiquitous penetration of mobile networks and software-defined networking technology enable us to sense, predict and control the physical world using information technology - the so-called Internet of Things (IoT)[1]. Consequently, business models and processes have been redesigned across a broad range of industries where objects are connected over the Web for communication with other objects on the Web, leading to the so-called Web of Things (WoT) on top of the IoT.

Pervasive connectivity, smart personal devices, for example in our homes, and demand for data testify to a WoT that will continue to grow. New devices are being developed and are becoming cheaper, making their integration into everyday objects ever more feasible, and as people buy into WoT technology, economies of scale lend themselves to the creation of ever more data-centric businesses.

The capabilities of these networks of devices presents us with several new and complex challenges that need to be solved before the Web of Things can deliver its promised potential. While there are, for example, some industry initiatives to achieve interoperability between smart home devices on the communication layer, including a recent collaboration between Google, Amazon, and Apple[2] to build a specific set of IP-based networking technologies for device certification, the data that these devices generate on the Web is not described uniformly. However, without connecting the data and its semantics that is generated by potentially billions of devices, the users of the WoT will end up in silos of information that require different applications to access and use it. A description of the capabilities of these devices and its context using semantic technologies may help in deciding how to communicate with the device and manage the data that is produced or the actions that can be performed.

At the data level this problem can be solved using an ontology-based approach. Gruber [24] introduced ontologies to Computer Science as an "*explicit specification of a conceptualization*" consisting of a "*set of objects, and the describable relationships among*

*Corresponding author. E-mail: maxime.lefrancois@emse.fr.
[1] also referred to as Cyber-Physical Systems (CPS)
[2] https://zigbeealliance.org/news_and_articles/connectedhomeip/

them" represented in a declarative formalism. Ontologies have consequently proven to be a very useful tool for semantic interoperability between parties that exchange data. With the adoption of the Linked Data movement, as proposed by Berners-Lee [5], machine-readable connections between data expressed in RDF in combination with ontologies have since seen a large uptake on the Web. One such ontology backed by a consortia of search engine providers including Google, schema.org, has seen particular success on the Web and has been used for semantic annotations on 31.3% of all websites [26] already back in 2015. It is now responsible for a large portion of the currently available machine-readable data on the Web.

There already exists a significant amount of research focusing on applying the RDF data model and OWL ontologies in different WoT scenarios, from home automation to Industry 4.0, by showing how this approach can be applied to ease integration of diverse data sources [39, 55]. Ontologies and vocabularies such as the Semantic Sensor Network Ontology (SSN) [29] have been adopted in a number of research projects [7, 63, 69]. Although the ontology-based approach in WoT has received significant interest and adoption in research projects, it still lacks similar levels of adoption to schema.org on the Web or adoption in industry, more generally [38].

This lack of adoption can be attributed to several challenges, the following three of which we consider as the major open challenges.

1. Maturity and Coverage of WoT Ontologies
2. Semantics in the Edge
3. Distributed and embedded reasoning

In the following section, we will detail and discuss these three challenges in more detail.

2. Challenges in deploying a Web of Things

2.1. Maturity and Coverage of WoT Ontologies

Different standardisation bodies work towards developing data models and ontologies for the Internet and Web of Things [18, 28, 34, 35]: the World Wide Web Consortium (W3C), the Open Geospatial Consortium (OGC), the Internet Engineering Task Force (IETF), the European Telecommunication Standards Institute (ETSI), the Open Connectivity Foundation (OCF), the IPSO Alliance, and the Open Mobile Alliance, among them. Many aspects of the data being generated in the WoT need to be described semantically and the standardisation bodies sometimes adopt conceptually different modelling perspectives. The diagram in Figure 1 shows these aspects and presents the current state-of-the-art in ontologies available to describe those.

Real world setting A central aspect in modeling WoT applications is the description of the real world setting in which the Things/Sensors/Actuators are deployed to observe or act on some features of interest. Many ontologies [4, 12, 29, 33, 61] include patterns to describe such settings. The OGC and W3C joint standard on the Semantic Sensor Network ontology (SOSA/SSN) [28], describing networks of sensors and actuators, their capabilities, their features of interest, and their individual observations or actuations, serves as a core or as a source of inspiration to many of these ontologies, ensuring some form of interoperability between them. The ETSI SmartM2M technical committee, develops the Smart Applications REFerence (SAREF) Ontology [12][3], to describe devices and their functions. SAREF is aligned with the oneM2M [52] base ontology that describes communication devices and the messages they exchange, for syntactic and semantic interoperability with external systems.

Extensions of SAREF and SSN have been developed for specific domains, such as CASO for the Agriculture [51], EEPSA for Buildings [23], and the SAREF4*ABCD* series of SAREF extensions[4] (e.g., [13, 56]). The topological organization of features of interest, which is the focus of [45, 46] is often an important aspect as properties of related features of interest may be inter-dependent. Specializations such as BOT [57] are defined for specific domains.

In addition, SSN has a separate module, called SSN System, to model capabilities and operating/survival ranges of systems/things. Sagar et al. [60] discussed some remaining modeling issues of SSN in this regard and proposed the S3N[5] extension focusing on modeling reconfiguration capabilities of sensors and actuators. It is a continuous challenge to make these different initiatives progressively converge.

Property qualification or quantification Sensors and actuators are deployed to observe and act on specific properties of features of interest. SOSA/SSN and SAREF both have the modeling of features of inter-

[3]SAREF ontology - https://saref.etsi.org/
[4]SAREF extensions - https://saref.etsi.org/extensions.html
[5]S3N ontology - https://w3id.org/s3n/

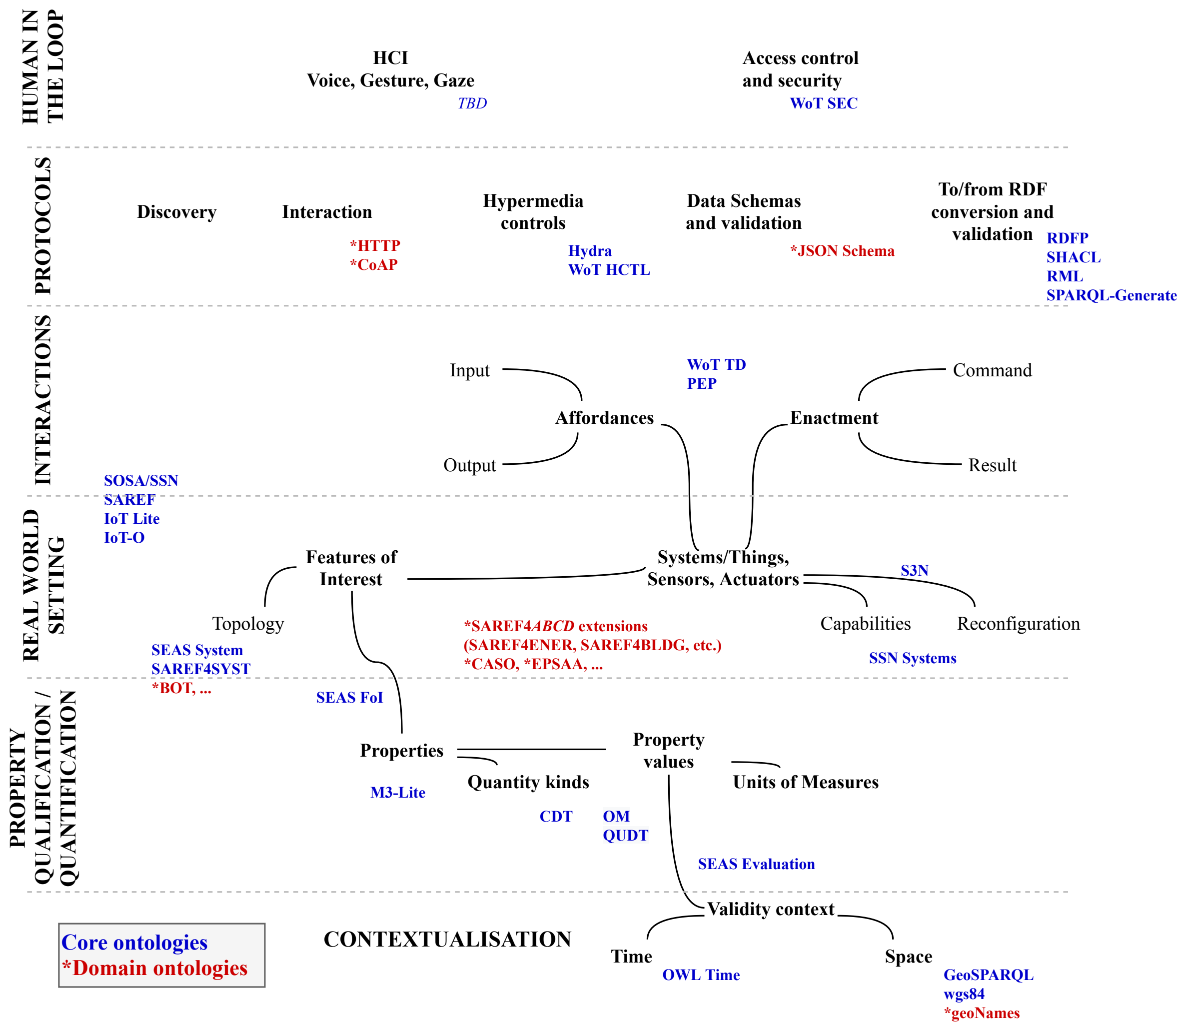

Fig. 1. Data aspects that require semantics in WoT and existing established ontologies

ests and their properties in their core, but do not prescribe whether a property can be reused across features of interest, or should belong to a specific feature of interest[6], which may lead to interoperability issues across datasets and ontologies. A more strict axiomatization was proposed in SEAS [46]. The work in QUDT [31] on a list of quantity kinds can also help in generically defining properties of features of interest in different domains, e.g. the ontology defines, among many other quantity kinds, a `quantitykind:AreaTemperature` that can be used in many contexts. Taxonomies of properties and sensors are defined in domain ontologies such as M3-Lite [1], which still need to be adapted to the new version of SSN.

Being able to describe quantity values and their units is a requirement that is almost ubiquitous in any WoT domain. Different ontologies have been developed to describe units, their relations, and quantities with their values. A recent survey [37] compares and evaluates eight well known ontologies for units of measurements, among which QUDT [31], OM [58] and the Units Ontology [22] are the most widely used. The survey also reports on the Wikidata corpus [68] that at the time of research contained over 4.4k measurement units and 4.1k non-prefixed units. While these ontologies are comprehensive in respect

[6]See Footnote 10 in [28] and https://www.w3.org/TR/vocab-ssn/#SSNProperty-instances for details on this point.

to modelling units of measurements and their relations, a comprehensive model of systems of quantity kinds is still under development, with the QUDT ontology leading the way (as mentioned above). An alternative approach relying on RDF 1.1 Datatypes is proposed by Lefrançois and Zimmermann [48], and allows for more concise representation of quantity values and queries.

Contextualization of value assignment to properties The OWL time ontology [10], the WGS84[7], or the GeoSPARQL [54] vocabulary, can be used to describe when and where an actuation or observation is made or valid, or when and where a property has a certain value. These ontologies may also be used to model spatial and temporal properties of devices and features of interests. GeoNames also provides a dataset of eleven million placenames that can be used for annotating locations using human-understandable labels. While OWL time fully covers the temporal requirements in WoT applications, there are still some unsolved issues in modelling spatial aspects [67], i.e. around indoor location relations, authoritative geometrical descriptions of place boundaries and relations that define qualitative assertions based on human perceptions to relate places that are deemed to be the same. However, these are relatively minor issues for modelling spatial properties in WoT applications.

Device functionality and their APIs A recent and important core ontology to describe thing affordances in terms of properties, actions, and events, is the Thing description [35] developed by the W3C Web of Things working group. However, the Thing description does not model how the enactment of these affordances is to be modeled. This aspect is covered by SOSA/SSN that describe sensors that implement procedures and make observation. Parallel to this, they describe actuators that implement procedures and make actuations. The PEP [46] ontology defines an ontology pattern as a generalization of these two parallel conceptual models. However, the aforementioned models are mainly concerned with the flow of information between devices, but there is little integration of device types (sensors, actuators, gateways) themselves. They need to be considered as more than raw data producers, but the data input and output of these devices and the process to execute them needs to undergo harmonisation through ontological models. This has not yet happened other than in some subdomains.

[7]W3C Basic Geo ontology - https://www.w3.org/2003/01/geo/

Events and Processes While the PROV Ontology (PROV-O) [42] is an established W3C standard that provides a set of classes, properties, and restrictions that can be used to represent events and activities that happened, i.e. to document a workflow log, an established process ontology that can be used to describe the execution behaviour of complex functions of WoT devices is missing. PSL [25] and m3po [27] were early examples of process ontologies that could theoretically be used for WoT devices, but neither are described in OWL nor have they found use outside of academia. The WiLD ontology [36] proposed as an execution model for the Linked-Data Fu system is the closest to a process model to describe the execution behaviour of WoT devices, but it has yet to be used in actual implementations and it also lacks a mapping to the ontologies mentioned above for describing device functionality and their APIs, in particular the WoT Thing description.

WoT protocols To be part of the WoT, things, sensors, and actuators, need to be exposed on the Internet and reachable using Web protocols. This category focuses on ontologies that bind the thing's affordances and their enactments to Web protocols. The architecture paradigm for WoT applications is intended to be stateless due to Thing constraints, therefore ontologies such as Hydra [40] are a good fit. More than the classical REST level 2 that is offered by Hydra, the WoT HCTL ontology[8] intends to further describe hypermedia controls (REST level 3), in the form of links and forms. Links are a transposition to RDF of the IETF RFC 8288 Web Links. The WoT HCTL is developed in parallel to a specification effort of the IRTF T2T group called Constrained RESTful Application Language (CoRAL)[9].

A missing link between the description of affordances and the actual messages that will be sent as commands and received as results, is the description of the data model for these messages. The WoT JSON Schema ontology[10] can be used with WoT TD to specify a data schema for these messages if they are using JSON. It remains a challenge to bring semantics in the edge such that WoT devices consume and produce RDF. Section 2.2 provides a general overview of this challenge.

[8]WoT HCTL - https://www.w3.org/2019/wot/hypermedia
[9]CoRAL - https://tools.ietf.org/html/draft-ietf-core-coral-03
[10]JSON Schema ontology - https://www.w3.org/2019/wot/json-schema

Human in the loop Voice and gesture-controlled interfaces are becoming increasingly popular in the WoT. In particular, in automobiles, smart homes, computer games and Augmented Reality (AR) / Virtual Reality (VR) applications, voice, gestures and sometimes even gaze has become prevalent due to its accessibility to everyone. Designers, producers, and vendors integrating gesture interfaces into their products have also increased in numbers, giving rise to a greater variation of interaction models in utilizing them. However, different modalities that are used to interact with smart environments have not yet been formalized in vocabularies and ontologies, in particular models that formally describe voice commands, gestures and gaze interactions and how they relate to affordances of WoT devices are missing.

Other considerations are of utmost importance for the WoT as it bridges the Web and the real-world where humans need to protect their privacy and integrity. Ontologies to describe access control and security of WoT application will be important for a successful deployment of the semantic WoT. The WoT SEC ontology[11] under development is a notable initiative in this line.

2.2. Semantics in the Edge

While RDF has proven to be an effective data model for interoperability on the application layer, its verbose serialisation formats (e.g. RDF/XML, NTriples, or Turtle) present a challenge on the presentation layer. Other than some approaches using the HDT [19] serialisation of RDF [30] or other binary representations of RDF [8], there has been little work and even fewer uptake in industry of providing WoT devices that consume and produce RDF.

Consequently, many data formats and data models exist and they compete with each other for adoption in devices in different WoT domains. Standardisation groups rather try to solve this problem by standardising data formats and service APIs [17, 21, 34]. Some work aim at tackling semantic interoperability despite the heterogeneity of data formats and service API specifications, i.e., across platforms.

The use of semantic Web technologies has been investigated to facilitate semantic interoperability among these platforms [20, 50, 65]. One challenge is to investigate how semantic interoperability can be obtained on the edge level, i.e. between devices directly, instead of between platforms. The work in Lefrançois [47] is a starting point to investigate how constrained devices that are not natively semantic Web enabled can still be interoperable with one another.

Figure 2a illustrates a typical Web service that consumes and outputs resource representations, that are octet streams typed with internet media types according to the Web architecture principles[12]. For some data formats such as JSON or XML, dedicated validation languages such as JSON Schema or XML Schema may be used. Then, the contents of the resource whose representation is given as input or output may be assumed to be an RDF graph. Adopting such an abstraction enables us to assume the service, potentially exposed by a constrained device, consumes and produces RDF. Many languages can be used to specify how an RDF graph can be generated out of octet streams (lifting), or the other way around (lowering). Finally languages such as SHACL or ShEX can be used to specify what form the content RDF graph has. New research challenges stem naturally from this vision. An example is that given a JSON Schema representation validation rule, and a lifting rule, to automatically compute the SHACL shape that the content should validate against.

Figure 2b illustrates a combination of two WoT services that are seemingly incompatible, but that as an abstraction generate and consume RDF, respectively. The output RDF graph, generated using a certain lifting rule, could then be lowered using the second thing's lowering rule. In this setting, the condition for the services to be composable is that the content validation rule of the first thing is more specific than the content validation rule of the second thing. However, SHACL shape containment has not been investigated yet.

2.3. Distributed and embedded reasoning in the Edge

As devices became powerful enough to offer storage and processing, new architectures appeared, based on edge computing. At the same time, it is now often a requirement that the final user is able to configure the intelligent environment. This poses several research questions: (i) how to embed reasoning in edge devices with various capacities; (ii) how to efficiently distribute reasoning tasks among available heterogeneous devices; and (iii) how to allow user to easily write rules for such devices.

[11] WoT SEC - https://www.w3.org/2019/wot/security

[12] https://www.w3.org/TR/webarch/#internet-media-type

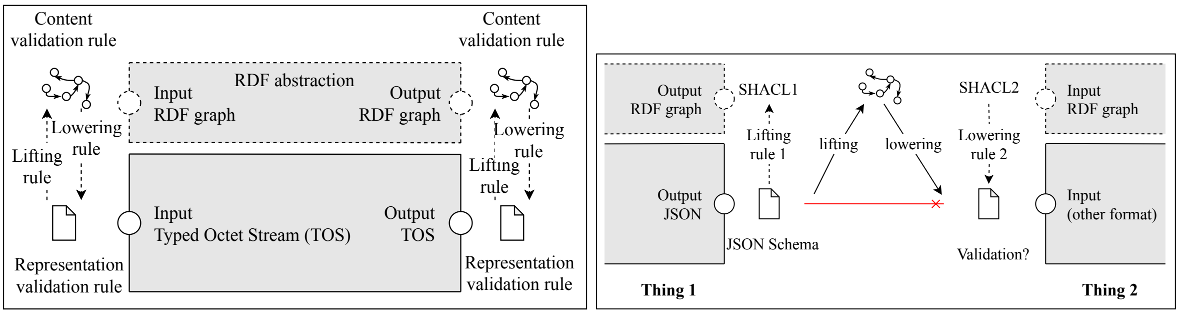

(a) Web service with RDF abstractions. (b) Heterogeneous things interacting on the WoT.

Fig. 2. Semantics in the Edge: using RDF abstractions for the content of Web resource.

How to embed reasoning in devices with various capacities Edge computing allows manipulating data close to sources, saving bandwidth, lowering latency and reducing communication needs. Slider [9] is an incremental reasoner optimised in memory and processing footprint. RDF4LED [41] is a lightweight RDF engine, which comprises of RDF storage and a SPARQL processor, for small query operations in lightweight edge devices. Devices on the WoT generate and consume highly dynamic data. The recent survey on stream reasoning by Dell'aglio et al. [14] discusses several open issues left to tackle in this field.

How to distribute reasoning tasks among devices Reasoning tasks can be distributed among heterogeneous devices: some are powerful computers, some are edge devices with various constraints. One challenge is to develop computationally-efficient reasoning strategies coping with such heterogeneities [62], and as close to the data sources as possible. One such approach is HyLAR that deploys incremental reasoning tasks on both, the server and the client [64]. Assuming that devices may reason with heterogeneous entailment regimes is also a starting point for a new investigation, for example, if a constrained edge node uses a poorer entailment regime than a more powerful gateway.

End-User Development In the scenario offered by such a complex network, the End-User Development (EUD) vision [49] aims at putting customization mechanisms in the hands of end users. Starting from iCAP [16], an early rule-based system for building context-aware applications, several works demonstrated the effective applicability of EUD techniques for the personalization of the functionality of smart devices and online services in different areas, including mobile environments [11] and smart home [6]. Users can take advantage of visual programming platforms such as IFTTT and Zapier to personalize the joint behaviors of their own connected entities, by adopting the trigger-action programming paradigm, i.e., they allow the definition of IF-THEN rules. Unfortunately, despite its wide adoption, the way it is implemented nowadays presents its own set of open issues: i) low-level of abstraction (current trigger-action programming platforms adopt highly technology-dependent representation models that work with well-known connected entities, previously associated to a specific user, only. Therefore, defining IF-THEN rules becomes a complex task for non-programmers [32]), ii) information overload (contemporary trigger-action programming platforms do not provide users with any discovery support [66], and the explosion of new smart devices and online service results in user interfaces with too much information), iii) and run-time problems (there is the need to provide users with instruments for understanding and debugging their IF-THEN rules, i.e., to avoid possible conflicts [6] and to assess a rules' correctness [15]).

3. Overview of the Special Issue

The focus of this special issue is to showcase novel approaches of applying semantic technologies to solve the problems of device and data integration mentioned above. We received nine submissions, covering a wide range of points of view related to these topics. The thorough peer-review process selected three of these submissions, of three different submission types (i.e. one full paper, one survey article and one linked dataset description), as mature enough to be published in the Semantic Web journal.

The three accepted papers also deal with three popular and important application domains of the WoT [59], *Smart Agriculture*, *Smart Building* and the *Industry of the Future*[13].

The paper on *Weather Data Publication on the LOD using SOSA/SSN Ontology*[14] by Catherine Roussey, Bernard Stephan, André Géraldine, and Boffety Daniel on *Smart Agriculture* is a typical application domain of WoT architectures, where semantically annotated weather/climate data [2, 44] and the monitoring of cultivated fields requires various sensors that push streaming data [63] that must be collected and reasoned upon to take decisions executed by actuators. This paper specifically presents an RDF dataset of meteorological measurements that have been obtained by a weather station at an experimental farm located in Montoldre, France and then be converted to Linked Open Data (LOD). The work reuses many of the established ontologies for WoT that we discuss in Section 2.1. At the core of the dataset sits the new SOSA [33] and SSN [29] ontologies, and extensions of SSN for meteorological sensors [43]. Further, to model the geospatial aspects of the sensors the OGC GeoSPARQL vocabulary [54] and the OWL Time ontology [10] are used, while for weather related measurements such as the temperature, precipitation, wind and solar radiation the QUDT ontology [31] is used.

The survey paper *Ontologies for Observations and Actuations in Buildings: A Survey*[15] by Iker Esnaola-Gonzalez, Jesús Bermúdez, Izaskun Fernandez, and Aitor Arnaiz, discusses another typical application domain, *smart buildings*, where added-value application services involve information from other verticals such as energy management, e-health, or ageing well. It reviews and compares existing ontologies in the IoT and building domain, using a set of competency questions extracted from a simple situation. Candidate ontologies are filtered based on base quality criteria such as whether the ontology is online, with metadata, documented, designed using principles, and used. Among the ten selected ontologies to model observations and actuations, many are well established and referenced in Figure 1, such as SOSA/SSN [29], SAREF [18], SEAS [46], IoT-O [61], IoT-Lite [4]. The authors then compare available ontologies for expressing time, location, and units of measurements and quantities, with extended conclusions with respect to Section 2.1. Finally, the authors review ten building domain ontologies, some of which are actively maintained by important consortia: ifcOWL [53] of BuildingSMART, BOT [57] of the W3C Linked Building Data community group, SAREF4BLDG [56] of ETSI SmartM2M, and Brick [3].

Industry of the Future or Industry 4.0 is a third interesting application domain, where devices, machines, production modules and products are comprised as Cyber-Physical Systems (CPS) that are autonomously exchanging information, triggering actions and controlling each other. Factories are developing into intelligent environments that enable dynamic re-engineering processes and the ability to respond flexibly to failures. In particular, business-specific knowledge must therefore be modeled as self-contained bundles, and inserted into the system at runtime when needed. To address this issues, the paper *EDR: A Generic Approach for the Distribution of Rule-Based Reasoning in a Cloud-Fog continuum*[16] by Nicolas Seydoux, Khalil Drira, Nathalie Hernandez, and Tierry Monteil, proposes an original architecture which exploits the complementarity of Cloud and Fog computing. In this model, reasoning rules are used to capture business level logic and are distributed across nodes and executed as close as possible to where the data is produced, in order to enable low-latency decision making. At the same time remote powerful Cloud computation resources are exploited in order to benefit from the Cloud stability and permanent availability. Moreover, as IoT networks are open and evolutive, the computation is dynamically distributed across Fog nodes according to the transformation of the network topology.

4. Conclusion and Future Directions

Contributions to this special issue have shown that ontologies, linked data, and reasoning, have a wide range of research directions on the Web of Things and can be applied to a wide range of application domains. However, further advances are needed to cover gaps in

[13] often also called Industry 4.0, a term originating in 2011 from a project in the high-tech strategy of the German government

[14] http://www.semantic-web-journal.net/content/weather-data-publication-lod-using-sosassn-ontology-0

[15] http://www.semantic-web-journal.net/content/ontologies-observations-and-actuations-buildings-survey-1

[16] http://www.semantic-web-journal.net/content/edr-generic-approach-distribution-rule-based-reasoning-cloud-fog-continuum

existing ontologies, bring semantics in the edge, and develop distributed semantic reasoning approaches.

5. Acknowledgements

We take this opportunity to sincerely thank the authors for their invaluable and inspiring contributions to this special issue. We are also grateful to all reviewers for reviewing the submissions and helping us publish an interesting special issue (non-anonymous reviewers only, including reviews on rejected papers and in alphabetical order): Jose María Alvarez-Rodríguez, Payam Barnaghi, Eva Blomqvist, Michel Böhms, Simon Cox, Aidan Hogan, Andreas Kamilaris, Sebastian Neumaier, Alessandro Oltramari, Antonio Piccinno, Ana Roxin, Antoine Zimmermann. Finally we thank the editors-in-chief of the Semantic Web Journal, Pascal Hitzler and Krzysztof Janowicz, for their continuous support and help and for making the Semantic Web Journal *Open Review* and *Open Access*.

References

[1] Rachit Agarwal, David Gomez Fernandez, Tarek Elsaleh, Amelie Gyrard, Jorge Lanza, Luis Sanchez, Nikolaos Georgantas, and Valerie Issarny. Unified iot ontology to enable interoperability and federation of testbeds. In *2016 IEEE 3rd World Forum on Internet of Things (WF-IoT)*, pages 70–75. IEEE, 2016. https://doi.org/10.1109/WF-IoT.2016.7845470.

[2] Ghislain Atemezing, Oscar Corcho, Daniel Garijo, José Mora, María Poveda-Villalón, Pablo Rozas, Daniel Vila-Suero, and Boris Villazón-Terrazas. Transforming meteorological data into linked data. *Semantic Web*, 4(3):285–290, July 2013. https://doi.org/10.3233/SW-120089.

[3] Bharathan Balaji, Arka Bhattacharya, Gabriel Fierro, Jingkun Gao, Joshua Gluck, Dezhi Hong, Aslak Johansen, Jason Koh, Joern Ploennigs, Yuvraj Agarwal, Mario Berges, David Culler, Rajesh E. Gupta, Mikkel Baun Kjærgaard, Mani B. Srivastava, and Kamin Whitehouse. Brick: Towards a Unified Metadata Schema For Buildings. In *Proceedings of the 3rd ACM International Conference on Systems for Energy-Efficient Built Environments*, pages 41–50. ACM, 2016. https://doi.org/10.1145/2993422.2993577.

[4] María Bermúdez-Edo, Tarek Elsaleh, Payam Barnaghi, and Kerry Taylor. IoT-Lite: A Lightweight Semantic Model for the Internet of Things. In *Proceedings of the International IEEE Conferences on Ubiquitous Intelligence Computing, Advanced and Trusted Computing, Scalable Computing and Communications, Cloud and Big Data Computing, Internet of People, and Smart World Congress*, pages 90–97, 2016. https://doi.org/10.1109/UIC-ATC-ScalCom-CBDCom-IoP-SmartWorld.2016.0035.

[5] Tim Berners-Lee. Linked data. Technical report, W3C, 2006. http://www.w3.org/DesignIssues/LinkedData.

[6] Danilo Caivano, Daniela Fogli, Rosa Lanzilotti, Antonio Piccinno, and Fabio Cassano. Supporting end users to control their smart home: design implications from a literature review and an empirical investigation. *Journal of Systems and Software*, 144:295–313, 2018. https://doi.org/10.1016/j.jss.2018.06.035.

[7] Jean-Paul Calbimonte, Hoyoung Jeung, Oscar Corcho, and Karl Aberer. Enabling query technologies for the semantic sensor web. *International Journal On Semantic Web and Information Systems (IJSWIS)*, 8(1):43–63, 2012. https://doi.org/10.4018/jswis.2012010103.

[8] Victor Charpenay, Sebastian Käbisch, and Harald Kosch. Towards a binary object notation for rdf. In *The Semantic Web*, pages 97–111, Cham, 2018. Springer. https://doi.org/10.1007/978-3-319-93417-4_7.

[9] Jules Chevalier, Julien Subercaze, Christophe Gravier, and Frédérique Laforest. Incremental and Directed Rule-Based Inference on RDFS. In *Database and Expert Systems Applications - 27th International Conference, DEXA 2016, Porto, Portugal, September 5-8, 2016, Proceedings, Part II*, pages 287–294. Springer, 2016. https://doi.org/10.1007/978-3-319-44406-2_22.

[10] Simon J D Cox and Chris Little. Time ontology in OWL. W3C Recommendation, W3C, October 2017. https://www.w3.org/TR/owl-time/.

[11] José Danado and Fabio Paternò. A prototype for EUD in touch-based mobile devices. In *2012 IEEE Symposium on Visual Languages and Human-Centric Computing (VL/HCC)*, pages 83–86. IEEE, 2012. https://doi.org/10.1109/VLHCC.2012.6344487.

[12] Laura Daniele, Frank den Hartog, and Jasper Roes. Created in Close Interaction with the Industry: The Smart Appliances REFerence (SAREF) Ontology. In Roberta Cuel and Robert Young, editors, *Formal Ontologies Meet Industry*, Berlin Heidelberg, 2015. Springer. https://doi.org/10.1007/978-3-319-21545-7_9.

[13] Mike de Roode, Laura Daniele, Alba Fernández-Izquierdo, Maria Poveda, and Raúl García-Castro. SAREF4INMA: a SAREF extension for the Industry and Manufacturing domain. *Semantic Web*, 2020. To appear in 2020.

[14] Daniele Dell'aglio, Emanuele Della Valle, Frank Van Harmelen, and Abraham Bernstein. Stream reasoning: A survey and outlook. *Data Science*, 1(1-2):59–83, 2017. https://doi.org/10.3233/DS-170006.

[15] Giuseppe Desolda, Carmelo Ardito, and Maristella Matera. Empowering end users to customize their smart environments: model, composition paradigms, and domain-specific tools. *ACM Transactions on Computer-Human Interaction (TOCHI)*, 24(2):1–52, 2017. https://doi.org/10.1145/3057859.

[16] Anind K Dey, Timothy Sohn, Sara Streng, and Justin Kodama. iCAP: Interactive Prototyping of Context-Aware Applications. In *Proceedings of the IEEE International Conference on Pervasive Computing (PerCom)*, pages 254–271. Springer, 2006. https://doi.org/10.1007/11748625_16.

[17] ETSI. Machine-to-Machine communications (M2M); mIa, dIa and mId interfaces. Technical Specification 102 921 V1.3.1, ETSI, September 2014.

[18] ETSI. SmartM2M; Smart Applications; Reference Ontology and oneM2M Mapping. Technical Specification 103 264 V3.1.1, ETSI, February 2020.

[19] Javier D. Fernández, Miguel A. Martínez-Prieto, Claudio Gutiérrez, Axel Polleres, and Mario Arias. Binary RDF representation for publication and exchange (HDT). *Journal of Web Semantics*, 19:22 – 41, 2013. https://doi.org/10.1016/j.websem.2013.01.002.

[20] Giancarlo Fortino, Claudio Savaglio, Carlos E Palau, Jara Suarez de Puga, Maria Ganzha, Marcin Paprzycki, Miguel Montesinos, Antonio Liotta, and Miguel Llop. Towards Multi-layer Interoperability of Heterogeneous IoT Platforms: The INTER-IoT Approach. In *Integration, Interconnection, and Interoperability of IoT Systems*, pages 199–232. Springer, 2018. https://doi.org/10.1007/978-3-319-61300-0_3.

[21] Open Connectivity Foundation. oneIoTa, 2017. URL http://oneiota.org/.

[22] Georgios V. Gkoutos, Paul N. Schofield, and Robert Hoehndorf. The units ontology: a tool for integrating units of measurement in science. *Database*, October 2012. 10.1093/database/bas033.

[23] Iker Esnaola Gonzalez. *Semantic technologies for supporting KDD processes*. PhD thesis, Universidad del País Vasco-Euskal Herriko Unibertsitatea, 2019.

[24] Tom Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5:199–220, 1993. https://doi.org/10.1145/503124.503150.

[25] Michael Grüninger. *Ontology of the Process Specification Language*, pages 575–592. Springer Berlin Heidelberg, Berlin, Heidelberg, 2004. https://doi.org/10.1007/978-3-540-24750-0_29.

[26] R.V. Guha, Dan Brickley, and Steve Macbeth. Schema.org: Evolution of structured data on the web. *ACM Queue*, 13(9), 2015. https://doi.org/10.1145/2857274.2857276.

[27] Armin Haller, Eyal Oren, and Paavo Kotinurmi. m3po: An ontology to relate choreographies to workflow models. In *Proceedings of the IEEE International Conference on Services Computing (SCC 2006), 18-22 September 2006, Chicago, Illinois, USA*, pages 19–27. IEEE Computer Society, 2006. https://doi.org/10.1109/SCC.2006.65.

[28] Armin Haller, Krzysztof Janowicz, Simon J D Cox, Danh Le Phuoc, Kerry Taylor, and Maxime Lefrançois. Semantic Sensor Network Ontology. W3C Recommendation, W3C, October 19 2017. https://www.w3.org/TR/vocab-ssn/.

[29] Armin Haller, Krzysztof Janowicz, Simon J. D. Cox, Maxime Lefrançois, Kerry Taylor, Danh Le Phuoc, Joshua Lieberman, Raúl García-Castro, Rob Atkinson, and Claus Stadler. The modular SSN ontology: A joint W3C and OGC standard specifying the semantics of sensors, observations, sampling, and actuation. *Semantic Web*, 10(1):9–32, 2019. https://doi.org/10.3233/SW-180320.

[30] Henning Hasemann, Alexander Kröller, and Max Pagel. RDF provisioning for the Internet of Things. In *2012 3rd IEEE International Conference on the Internet of Things*, pages 143–150, 2012. https://doi.org/10.1109/IOT.2012.6402316.

[31] Ralph Hodgson, Daniel Mekonnen, David Price, Jack Hodges, James E. Masters, Simon Cox, and Steve Ray. Quantities, Units, Dimensions and Types (QUDT) Schema - Version 2.1.2. Technical report, qudt.org, april 2020. http://www.qudt.org/doc/2020/04/DOC_SCHEMA-QUDT-v2.1.2.html.

[32] Ting-Hao Kenneth Huang, Amos Azaria, Oscar J. Romero, and Jeffrey P. Bigham. InstructableCrowd: Creating IF-THEN Rules for Smartphones via Conversations with the Crowd. *Human Computation*, 6(1):113–146, 2019.

[33] Krzysztof Janowicz, Armin Haller, Simon J.D. Cox, Danh Le Phuoc, and Maxime Lefrançois. SOSA: A lightweight ontology for sensors, observations, samples, and actuators. *Journal of Web Semantics*, 56:1–10, May 2019. https://doi.org/10.1016/j.websem.2018.06.003.

[34] Cullen Jennings, Z Shelby, J Arkko, A Keranen, and Carsten Bormann. Sensor Measurement Lists (SenML). RFC 8428, IETF, August 2018.

[35] Sebastian Kaebisch, Takuki Kamiya, Michael McCool, Victor Charpenay, and Matthias Kovatsch. Web of Things (WoT) Thing Description. W3C recommendation, W3C, April 2020.

[36] Tobias Käfer and Andreas Harth. Specifying, Monitoring, and Executing Workflows in Linked Data Environments. In *Proceedings of the 17th International Semantic Web Conference*, pages 424–440, Cham, 2018. Springer. https://doi.org/10.1007/978-3-030-00671-6_25.

[37] Jan Martin Keil and Sirko Schindler. Comparison and evaluation of ontologies for units of measurement. *Semantic Web*, 10:1–19, 08 2018. https://doi.org/10.3233/SW-180310.

[38] Maxim Kolchin, Nikolay Klimov, Alexey Andreev, Ivan Shilin, Daniil Garayzuev, Dmitry Mouromtsev, and Danil Zakoldaev. Ontologies for web of things: A pragmatic review. In *International Conference on Knowledge Engineering and the Semantic Web*, pages 102–116. Springer, 2015. https://doi.org/10.1007/978-3-319-24543-0_8.

[39] Konstantinos Kotis and Artem Katasonov. Semantic interoperability on the web of things: The semantic smart gateway framework. In *2012 Sixth International Conference on Complex, Intelligent, and Software Intensive Systems*, pages 630–635. IEEE, 2012. https://doi.org/10.1109/CISIS.2012.200.

[40] Markus Lanthaler. Hydra Core Vocabulary, A Vocabulary for Hypermedia-Driven Web APIs. Unofficial draft, Hydra W3C Community Group, October 2019.

[41] Anh Le-Tuan, Conor Hayes, Marcin Wylot, and Danh Le-Phuoc. RDF4Led: An RDF Engine for Lightweight Edge Devices. In *Proceedings of the 8th International Conference on the Internet of Things*, IOT '18, New York, NY, USA, 2018. Association for Computing Machinery. https://doi.org/10.1145/3277593.3277600.

[42] Timothy Lebo, Satya Sahoo, and Deborah McGuinness. PROV-O: The PROV Ontology. Technical report, W3C Recommendation 30 April, 2013.

[43] Laurent Lefort. Ontology for Meteorological sensors. Technical report, CSIRO, 2011. http://www.w3.org/2005/Incubator/ssn/ssnx/meteo/aws.

[44] Laurent Lefort, Armin Haller, Kerry Taylor, Geoffrey Squire, Peter Taylor, Dale Percival, and Andrew Woolf. The ACORN-SAT linked climate dataset. *Semantic Web*, 8(6):959–967, 2017. https://doi.org/10.3233/SW-160241.

[45] Maxime Lefrançois. SAREF consolidation with new reference ontology patterns, based on the experience from the SEAS project. TS 103 548 V1.1.2, ETSI.

[46] Maxime Lefrançois. Planned ETSI SAREF Extensions based on the W3C&OGC SOSA/SSN-compatible SEAS Ontology Pattern. In *Joint Proceedings of SEMANTiCS 2017 Workshops*. CEUR, 2017.

[47] Maxime Lefrançois. RDF presentation and correct content conveyance for legacy services and the web of things. In *Proceedings of the 8th International Conference on the Internet of Things, IOT 2018*, volume 43, pages 1–8. ACM, 2018. https://doi.org/10.1145/3277593.3277618.

[48] Maxime Lefrançois and Antoine Zimmermann. The unified code for units of measure in rdf: cdt: ucum and other ucum datatypes. In *European Semantic Web Conference*, pages 196–201. Springer, 2018. https://doi.org/10.1007/978-3-319-98192-5_37.

[49] Henry Lieberman, Fabio Paternò, Markus Klann, and Volker Wulf. End-user development: An emerging paradigm. In *End user development*, pages 1–8. Springer, 2006. https://doi.org/10.1007/1-4020-5386-X_1.

[50] Aref Meddeb. Internet of things standards: who stands out from the crowd? *IEEE Communications Magazine*, 54(7):40–47, 2016. https://doi.org/10.1109/MCOM.2016.7514162.

[51] Quang-Duy Nguyen, Catherine Roussey, María Poveda-Villalón, Christophe de Vaulx, and Jean-Pierre Chanet. Development Experience of a Context-Aware System for Smart Irrigation Using CASO and IRRIG Ontologies. *Applied Sciences*, 10(5):1803, 2020. https://doi.org/10.3390/app10051803.

[52] oneM2M. oneM2M Base Ontology. Technical Specification 0012 v3.6.0, oneM2M, 2018.

[53] Pieter Pauwels and Walter Terkaj. EXPRESS to OWL for construction industry: Towards a recommendable and usable ifcOWL ontology. *Automation in Construction*, 63:100–133, 2016. https://doi.org/10.1016/j.autcon.2015.12.003.

[54] Matthew Perry and John Herring. OGC GeoSPARQL - A Geographic Query Language for RDF Data. Ogc implementation standard ogc 11-052r4, OGC, Sept 2012. http://www.opengis.net/doc/IS/geosparql/1.0.

[55] Dennis Pfisterer, Kay Romer, Daniel Bimschas, Oliver Kleine, Richard Mietz, Cuong Truong, Henning Hasemann, Alexander Kröller, Max Pagel, Manfred Hauswirth, et al. Spitfire: toward a semantic web of things. *IEEE Communications Magazine*, 49(11):40–48, 2011. https://doi.org/10.1109/MCOM.2011.6069708.

[56] Marıa Poveda-Villalón and Raúl Garcıa-Castro. Extending the SAREF ontology for building devices and topology. In *Proceedings of the 6th Linked Data in Architecture and Construction Workshop, London, United Kingdom, June 19-21, 2018*, pages 16–23. CEUR, 2018.

[57] Mads Holten Rasmussen, Maxime Lefrançois, Georg Ferdinand Schneider, and Pieter Pauwels. BOT: the Building Topology Ontology of the W3C Linked Building Data Group. *Semantic Web*, 2020. To appear in 2020.

[58] Hajo Rijgersberg, Mark van Assem, and Jan Top. Ontology of units of measure and related concepts. *Semantic Web*, 4(1):3–13, 2013. https://doi.org/10.3233/SW-2012-0069.

[59] Marta Sabou, Stefan Biffl, Alfred Einfalt, Lukas Krammer, Wolfgang Kastner, and Fajar Juang Ekaputra. Semantics for Cyber-Physical Systems: A cross-domain perspective. *Semantic Web*, 11:115–124, 2020. https://doi.org/10.3233/SW-190381.

[60] Samya Sagar, Maxime Lefrançois, Issam Rebai, Khemaja Maha, Serge Garlatti, Jamel Feki, and Lionel Médini. Modeling smart sensors on top of SOSA/SSN and WoT TD with the semantic smart sensor network (S3N) modular ontology. In Elena Demidova, Amrapali J. Zaveri, and Elena Simperl, editors, *Emerging Topics in Semantic Technologies, ISWC 2018 Satellite Events*, Berlin, 2018. AKA Verlag.

[61] Nicolas Seydoux, Khalil Drira, Nathalie Hernandez, and Thierry Monteil. IoT-O, a Core-Domain IoT Ontology to Represent Connected Devices Networks. In Eva Blomqvist, Paolo Ciancarini, Francesco Poggi, and Fabio Vitali, editors, *Knowledge Engineering and Knowledge Management*, pages 561–576, Cham, 2016. Springer. https://doi.org/10.1007/978-3-319-49004-5_36.

[62] Xiang Su, Ekaterina Gilman, Peter Wetz, Jukka Riekki, Yifei Zuo, and Teemu Leppänen. Stream Reasoning for the Internet of Things: Challenges and Gap Analysis. In *Proceedings of the 6th International Conference on Web Intelligence, Mining and Semantics (WIMS)*, pages 1–10. ACM, 2016. https://doi.org/10.1145/2912845.2912853.

[63] Kerry Taylor, Colin Griffith, Laurent Lefort, Raj Gaire, Michael Compton, Tim Wark, David Lamb, Greg Falzon, and Mark Trotter. Farming the web of things. *IEEE Intelligent Systems*, 28(6):12–19, 2013. https://doi.org/10.1109/MIS.2013.102.

[64] Mehdi Terdjimi, Lionel Médini, and Michael Mrissa. HyLAR+: Improving Hybrid Location-Agnostic Reasoning with Incremental Rule-based Update. In *Proceedings of the 25th International Conference on World Wide Web, WWW 2016, Montreal, Canada, April 11-15, 2016, Companion Volume*, pages 259–262, 2016. https://doi.org/10.1145/2872518.2890542.

[65] Aparna Saisree Thuluva, Darko Anicic, and Sebastian Rudolph. IoT Semantic Interoperability with Device Description Shapes. In *The Semantic Web: ESWC 2018 Satellite Events - ESWC 2018 Satellite Events Revised Selected Papers*, pages 409–422. Springer, 2018. https://doi.org/10.1007/978-3-319-98192-5_56.

[66] Blase Ur, Melwyn Pak Yong Ho, Stephen Brawner, Jiyun Lee, Sarah Mennicken, Noah Picard, Diane Schulze, and Michael L Littman. Trigger-action programming in the wild: An analysis of 200,000 ifttt recipes. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, pages 3227–3231, 2016. https://doi.org/10.1145/2858036.2858556.

[67] Linda van den Brink, Payam Barnaghi, Jeremy Tandy, Ghislain Atemezing, Rob Atkinson, Byron Cochrane, Yasmin Fathy, Raúl García Castro, Armin Haller, Andreas Harth, Krzysztof Janowicz, Şefki Kolozali, Bart van Leeuwen, Maxime Lefrançois, Josh Lieberman, Andrea Perego, Danh Le-Phuoc, Bill Roberts, Kerry Taylor, and Raphäel Troncy. Best practices for publishing, retrieving, and using spatial data on the web. *Semantic Web*, 10(1):95–114, 2019. https://doi.org/10.3233/SW-180305.

[68] Denny Vrandecic and Markus Krötzsch. Wikidata: a free collaborative knowledgebase. *Communication of the ACM*, 57 (10):78–85, 2014. https://doi.org/10.1145/2629489.

[69] Peter Wetz, Tuan-Dat Trinh, Ba-Lam Do, Amin Anjomshoaa, Elmar Kiesling, and A Min Tjoa. Towards an environmental information system for semantic stream data. In *Proceedings of the 28th Conference on Environmental Informatics - Informatics for Environmental Protection, Sustainable Development and Risk Management*, pages 637–644, 2014.