Semantic Interoperability for the Web of Things[1]

Editor

Murdock, Paul	Landis+Gyr

Contributors

Bassbouss, Louay	Fraunhofer FOKUS
Bauer, Martin	NEC
Ben Alaya, Mahdi	Sensinov
Bhowmik, Rajdeep	Binghamton University
Brett, Patrica	Honeywell
Chakraborty, Rabindra	Senslytics
Dadas, Mohammed	Orange
Davies, John	BT plc
Diab, Wael	Huawei
Drira, Khalil	LAAS CNRS
Eastham, Bryant	OpenDOF
El Kaed, Charbel	Schneider Electric
Elloumi, Omar	Nokia
Girod-Genet, Marc	TELECOM SudParis
Hernandez, Nathalie	IRIT
Hoffmeister, Michael	FESTO
Jiménez, Jaime	Ericson
Kanti Datta, Soumya	Eurecom
Khan, Imran	Schneider Electric
Kim, Dongjoo	LG
Kraft, Andreas	Deutsche Telekom
Logvinov, Oleg	IoTecha
Longstreth, Terry	Comcast
Martigne, Patrica	Orange
Mladin, Catalina	InterDigital
Monteil, Thierry	LAAS CNRS
Murdock, Paul	Landis+Gyr
Nappey, Philippé	Schneider Electric
Raggett, Dave	W3C
Roes, Jasper	TNO
Serrano, Martin	Insight Centre for Data Analytics
Seydoux, Nicolas	LAAS CNRS
Simmon, Eric	NIST
Subramaniam, Ravi	Intel
Swetina, Joerg	NEC
Underwood, Mark	Krypton Brothers
Wang, Chonggang	InterDigital
Whitehead, Cliff	Rockwell Automation
Zhang, Yongjing	Huawei

Background

This paper is co-authored by an informal group of experts from a broad range of backgrounds all of whom are active in standards groups, consortia and/or alliances in the Internet of Things (IoT) space.

The ambition is to create mindshare on approaches to semantic interoperability and to actively encourage consensus building on what the co-authors regard as a key technical issue.

The paper

- Considers the value associated with interoperability in the IoT context and suggests that building mindshare across the industry on semantic approaches is one of the keys to unlocking that potential
- Introduces foundational interoperability concepts and provides a discussion of metadata, the rationale for sharing metadata, and the requirements for ontologies
- Introduces ontologies and discusses their specific relevance to interoperability and significance in the IoT context
- Provides examples of modular ontologies, overviews semantic annotation and tagging, and highlights strategies for scaling
- Draws conclusions and makes a number of recommendations

Introduction

The Internet of Things (IoT) generates expectations that smart devices can discover their context and build collaborations with other smart devices and services to create value. For example, smart devices in the home should be able to discover each other and to work together to both enhance the comfort and security of the home owner and to improve the efficiency of the home. When driving into the city, a smart car should be able to interact with city services to identify and reserve a parking place and should be able to collaborate with a personal smart phone to facilitate payment.

The expectation that ad-hoc networks of smart devices and services can be constantly formed and re-formed to manifest transient value systems is driving the need for broad agreement on how such devices interoperate and understand each other.

Discovery, understanding and collaboration at this level requires more than just an ability to interface and to exchange data. Whereas interoperability is "*the ability of two or more systems or components to exchange data and use information*" [1] , semantic interoperability "*means enabling different agents, services, and applications to exchange information, data and knowledge in a meaningful way, on and off the Web*" [2] .

Semantic interoperability is achieved when interacting systems attribute the same meaning to an exchanged piece of data, ensuring consistency of the data across systems regardless of individual data format. This consistency of meaning can be derived from pre-existing standards or agreements on the format and meaning of data or it can be derived in a dynamic way using shared vocabularies either in a schema form and/or in an ontology-driven approach. In this paper we will use the term "data-model based semantic interoperability" to refer to the former, and "ontology based semantic interoperability" to refer to the latter.

This paper considers semantic interoperability in the context of the Internet of Things (IoT). Note that we use "IoT" as an umbrella term for the range of emerging technologies which may differ in scope and reach[2], but which enable cross-domain innovation and drive the need for interoperability at a dynamic level.

[2] See [3] for a view on the range of IoT technologies and their scope

Semantic Interoperability as a Value Enabler

There are many analyst studies describing IoT as a broad concept spanning all application domains.

Beecham Research provided some early insight into the scope of domains covered with their "M2M Sector Map" [4] . A more recent study by McKinsey [5] considers the value potential of IoT in the context of domains broadly similar to those identified in the Beecham research.

The McKinsey study goes on to provide an estimation of the value that could be unlocked given interoperability across those domains – see Figure 1.

According to McKinsey:

Interoperability between IoT systems is critically important to capturing maximum value; on average, interoperability is required for 40 percent of potential value across IoT applications and by nearly 60 percent in some settings.

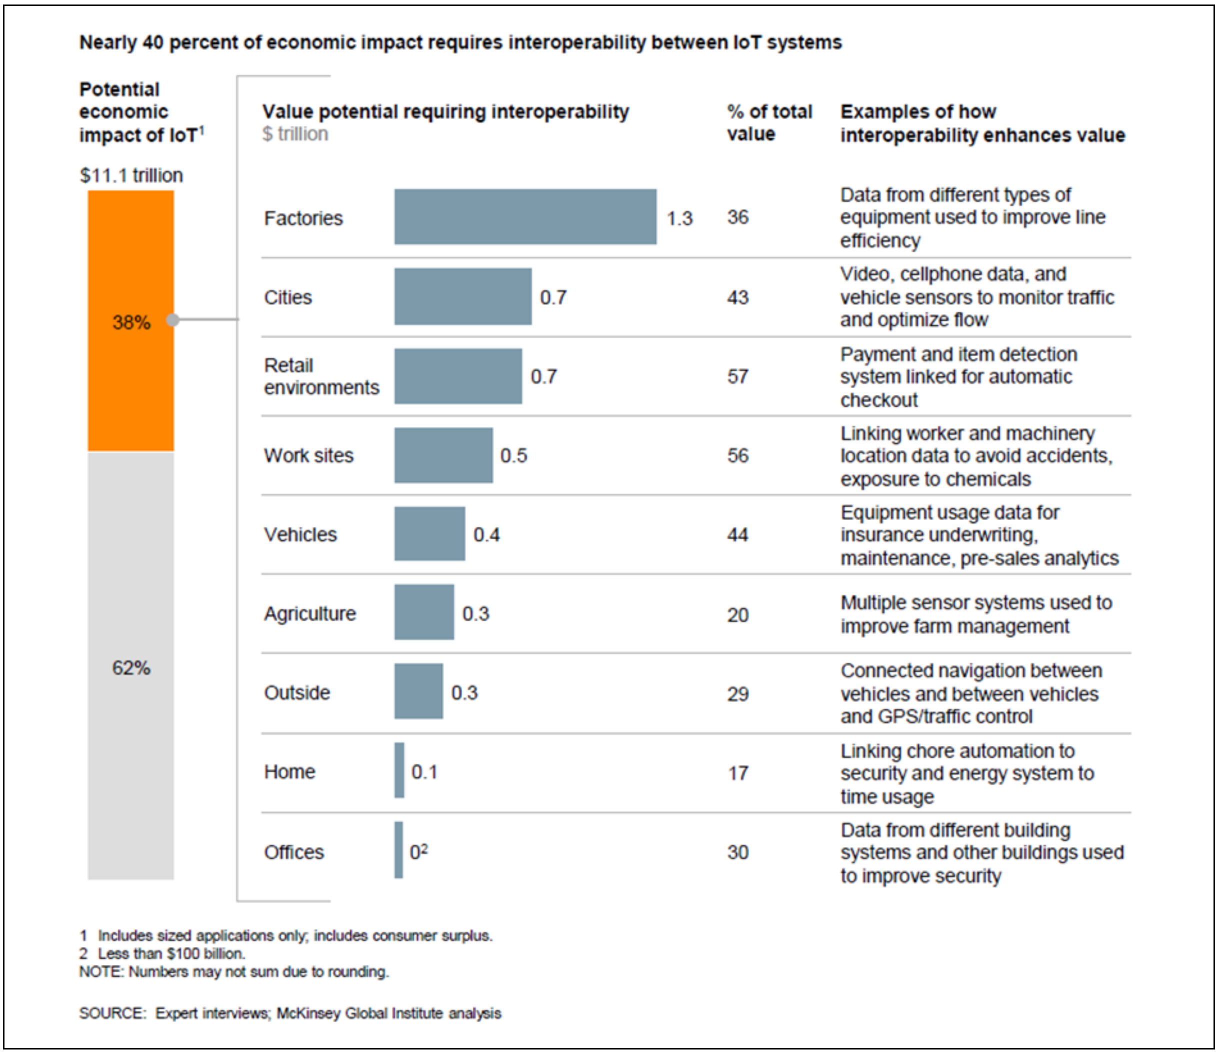

Figure 1[3] Value Potential Requiring Interoperability

We suggest that this McKinsey finding can be qualified further. Specifically, the full value potential can only be unlocked if interoperability is implemented in such a way that the dynamic nature of IoT is fully supported. A clear prerequisite is that ad-hoc, cross-domain systems of IoT elements need to

[3] “The internet of things: Mapping the value beyond the hype”, June 2015, McKinsey Global Institute, www.mckinsey.com.

be able to establish conversations and build understanding. Several examples of cross-domain use cases and a discussion on the need for and the value of semantics in sensor applications are described in [6] .

Hence, the position of this paper is that semantic interoperability is a key value enabler for IoT and that establishing a shared ontology based approach is critical for the development and exploitation of the technology.

Foundations of Semantic Interoperability

Metadata, an essential data reusability provider

What is the issue we want to solve?

Interoperability is commonly driven by the respective parties sharing *a priori* knowledge of some kind; for example, shared knowledge of an application programming interface (API) or shared knowledge of a set of database tables and related access rules. Key to these approaches is conformance with prior agreements and understandings.

These data-model based semantic interoperability approaches are the cornerstone of inter-operability in many enterprises and industrial contexts. One of the challenges, however, is that when new applications are introduced into the context, they also need prior knowledge of the interoperability schemes, the API specifications, and the meaning and use of database tables.

In the context of the IoT, however, there has to be a way to create an interoperability context which does not rely on prior knowledge. This is the issue we are trying to solve and metadata is a core part of the solution.

What is metadata?

Metadata is about describing the contents and context of data to facilitate discovery, understanding and (re)usability of that data. Hence the usual statement that metadata is data about data.

It is often the case, however, that the actual meaning of data can only be discovered by examining the software that generates and processes the data. This bundling obfuscates the semantics of the data with the result that third-party processors receiving data have no guidance on how the values should be interpreted and understood.

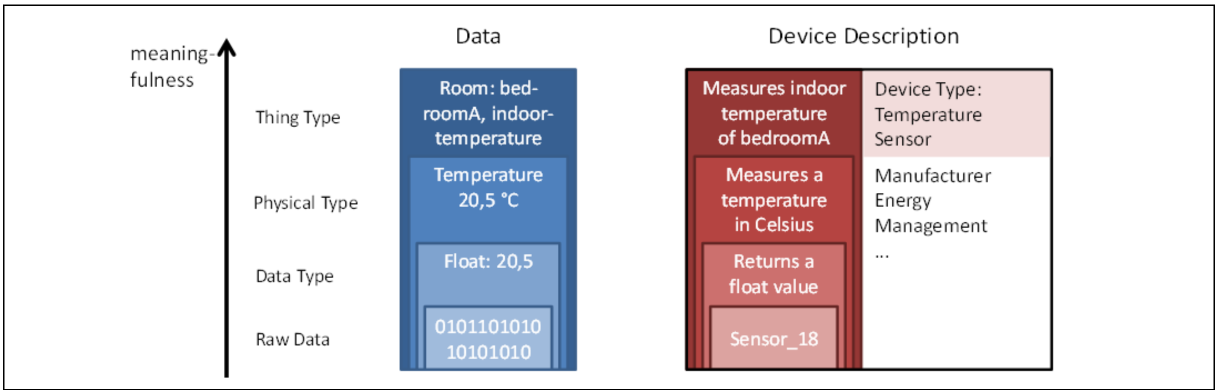

Figure 2 Meaningfulness of the data, increased with metadata

Metadata is about reducing the separation between semantics and values by ensuring that data is provided with context and description. This enables interpretation and understanding by subsequent processors and provides foundational support for interoperability and reusability.

Figure 2 [7] provides a view on the metadata associated with a temperature sensor. As can be seen, multiple levels of meaning can be inferred; this gives transparency to the context and supports subsequent design-time abstraction and modeling views.

Sharing metadata

The Linked Open Vocabularies[4] (LOV) community is driving a conversation around the creation and use of shared metadata (shared vocabularies).

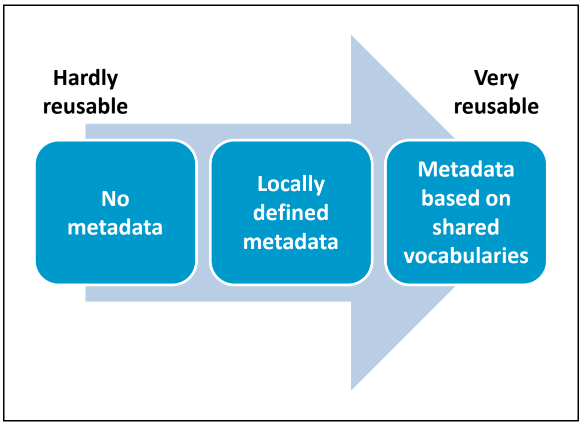

Figure 3 Shared Metadata

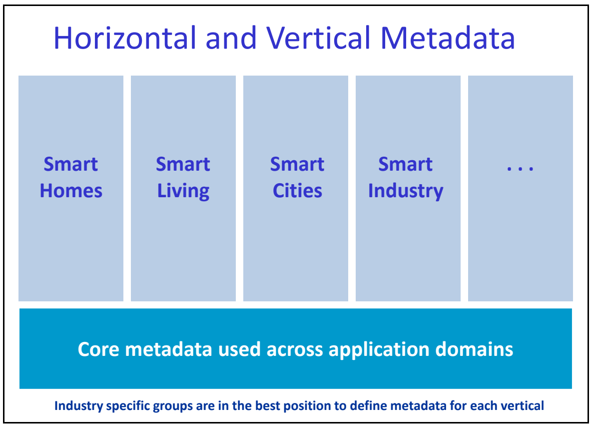

Figure 4 Metadata and Data Reusability

While locally defined and shared metadata will deliver value within a given domain, metadata which is published more widely will necessarily drive interoperability and reusability to a greater extent, as shown in Figure 3.

Programs such as the H2020 Large Scale Pilots [8] which address domain specific and cross-domain concerns, as shown in Figure 4, provide multiple opportunities for exercising and proving the strategic value of sharing metadata across a significant scope at scale.

Ontologies

What is an Ontology?

The LOV community is focusing on the curation of quality vocabularies across all domains. Taxonomies often build on such controlled vocabularies using parent-child relationships to describe the organization of terms within a specific domain. Ontologies extend this concept further to capture relationships capable of supporting richer operations and more advanced levels of reasoning.

Ontologies build on metadata to provide a representation of knowledge about a given domain and to provide a core resource for reasoning about a domain and a context. The Semantic Sensor Network (SSN) Ontology [9] is an example of an existing ontology which describes the capabilities and properties of sensors, the act of sensing and the resulting observations. Another example is the oneM2M Base Ontology [10] that constitutes a framework for specifying the semantics of data that are handled in oneM2M and to which domain specific ontologies can be mapped.

Figure 5 shows the key concepts and relations of the SSN ontology split by conceptual modules (dotted lines).

Figure 6 shows the core concepts of the oneM2M Base Ontology.

[4] http://lov.okfn.org/

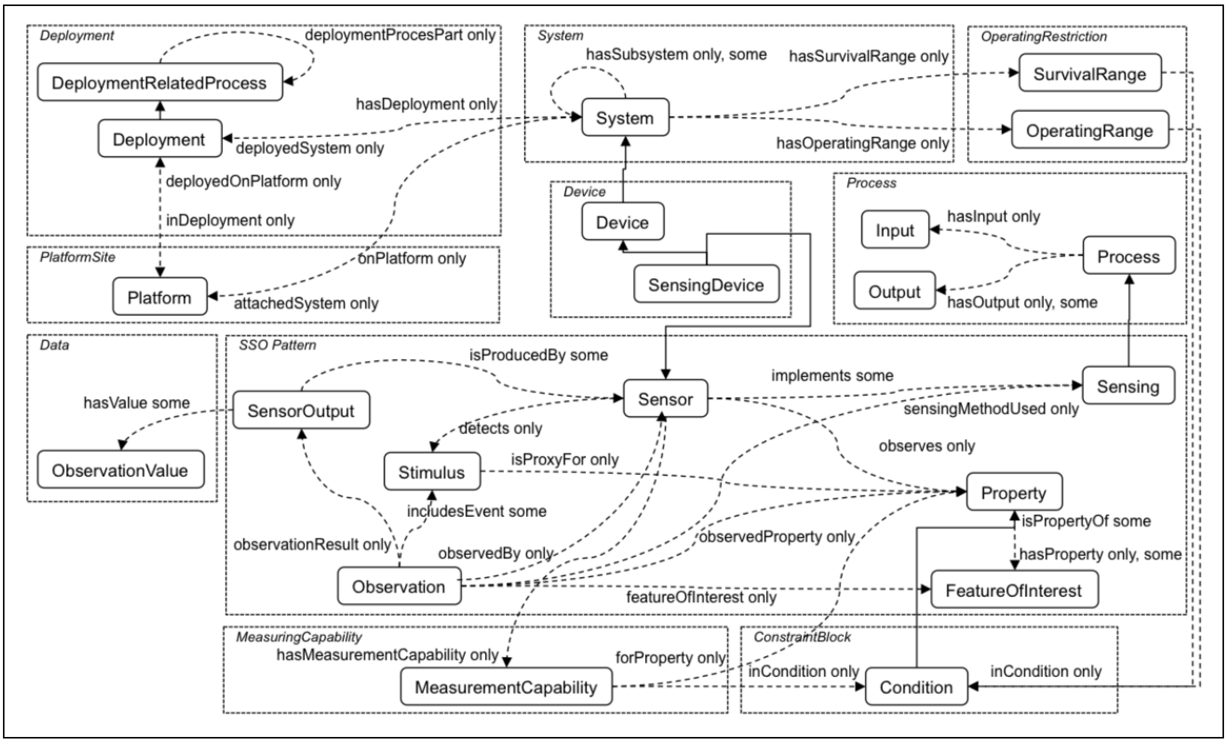

Figure 5 Semantic Sensor Network Ontology

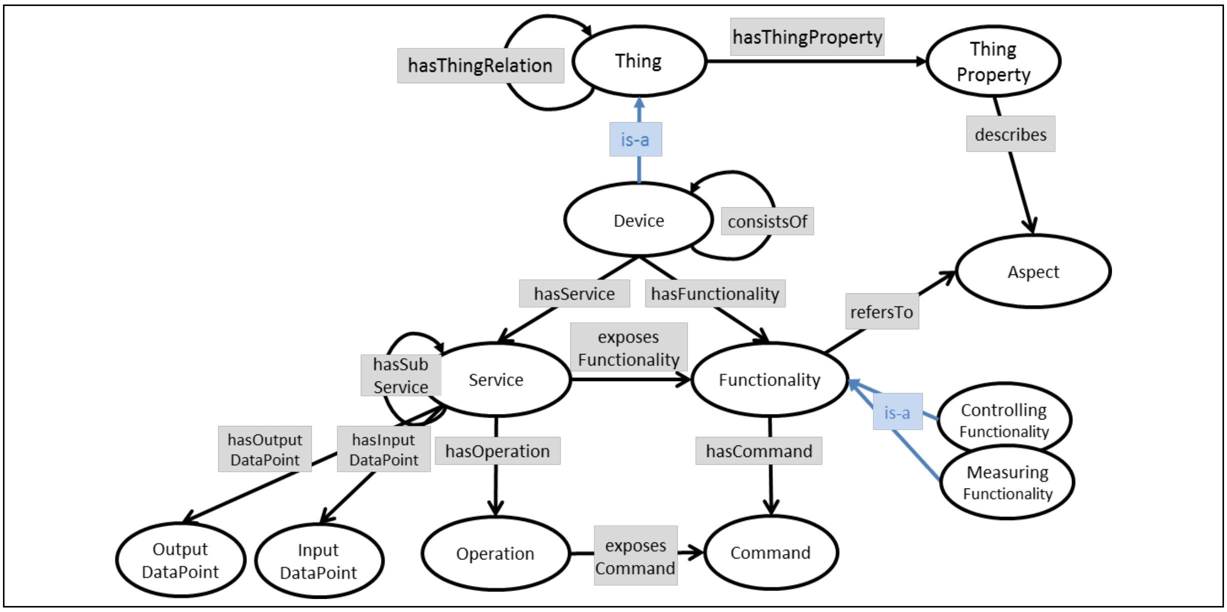

Figure 6 oneM2M Base Ontology

Ontologies and the IoT

Given the cross-domain nature of the IoT, there is a need both to capture and express knowledge shared across the verticals and to leverage linkages between domains.

As can be seen in Figure 5, the SSN ontology comprises ten conceptual modules relating to sensors. This modularity supports reuse of SSN concepts in other ontologies and, similarly, concepts from other ontologies can be included into solutions using the SSN ontology as required.

In Figure 6 the core aspects of the oneM2M Base Ontology are shown. The ontology provides a common, domain-independent basis to which existing domain-specific ontologies, e.g. SAREF [11] , can be mapped. IoT devices described according to the concepts of the Base Ontology, or derived from concepts in domain-specific ontologies, can be automatically mapped to a REST resource structure in oneM2M.

Modularity, reuse and linkage are key strategies for supporting the use of ontologies in building cross-domain IoT applications. These strategies, plus the need to educate the community in their existence and usage are discussed in [12] [13] .

The diversity of IoT domains will drive ontology development through vocabulary creation, extension, reuse, and retargeting; this motivates requirements for ontology management capabilities. The richness of semantic data models can be leveraged to automate management capabilities - such automation will be crucial for the continued operation of IoT ecosystems in which human intervention is expected to be minimal, ineffective and/or unavailable.

Ontologies and Modularity

Requirements for modularization are commonly driven by use-cases in which only parts of an existing ontology are needed, or in which constrained devices are unable to perform inference and reasoning on a full ontology. Modularization also eases some of the complexities around semantic data modelling and ontology design, integration, maintenance, and reuse [14] .

Modularization requires the partitioning of ontologies into independent sub-modules [15] [16] . Sub-modules are self-contained knowledge components that:

- Are loosely coupled
- Define their own set of core concepts and relations
- Are reusable
- Are linked to other module(s) with explicit relationship(s).

As a consequence of the loose coupling, modules can be designed, used, managed and updated in a stand-alone manner, with no impact on other modules. When modularizing ontologies, however, it's also important to avoid generating reasoning or querying complexities for future (module) unions.

Good examples of modular ontologies are the Smart BANs (Body Area Networks) and MyOntoSens ontologies [17] [18] . Within MyOntoSens, a Wireless Sensor Network (WSN) module is formed of clusters (Cluster module; BAN module for Smart BANs) that are composed of nodes (Node Module). A node is used for process (Process Module) and takes measurements (Measurements Module). The 'Measurements Module' is sufficiently light to be instantiated and stored within sensors, while the Process and Measurements modules full instantiation and inference/reasoning can actually only be performed within a more capable node, the cluster sink (or BAN hub). The full BAN ontology (including service level modules), is instantiated, inferred and processed within remote and distributed monitoring and control servers (e.g. hospital servers).

Ontology modularity can also be handled using a layered approach as shown in Figure 7. Although heterogeneity characterizes the landscape of devices and systems across domains, there are commonalties which can be abstracted out. Thus, ontologies can often be modularized into at least two layers: cross-domain ontologies and domain ontologies.

The cross-domain ontologies consist of concepts shared across domains and silos. For instance, a general protocol ontology can be used to classify the communication protocols along with information regarding the supported communication medium and range. Such general information can be used during diagnostic and maintenance operations. Similarly, there can be multiple cross-domain ontologies covering shared concepts related to quantities, units, topological relations, location, and usage.

The cross-domain ontologies capture the shared concepts across domains and constitute the building blocks of future extensions.

The domain ontologies relate to specific silos or verticals and often reference the cross-domain ontologies. For example, in Figure 7 the Buildings Ontology relies on both the Physical Quantities Ontology to express the measurements, and on the Localization Ontology to reference a site or a floor. Moreover, both cross-domain and domain ontologies can also rely on existing dictionaries (such as HayStack[5] from the building domain).

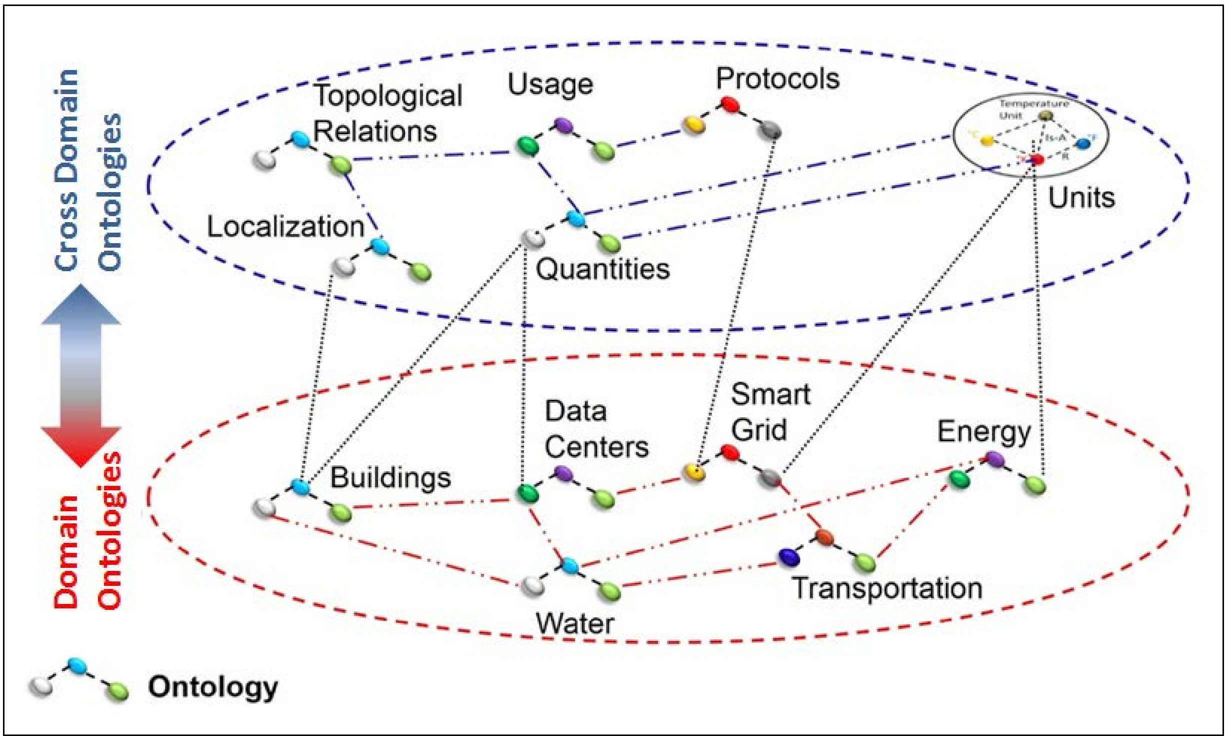

Figure 7 Multi-layered Ontologies

Consider the concepts of Current and Phase in the Buildings context:

- The Quantities Ontology (cross-domain) includes definitions for Current and Phase and also defines the hasPhase relation. Now assume that A is an instance of Phase.
- The Buildings Ontology (vertical domain) has a similar model and represents a current of phase A as IA. Thus, using the cross-domain ontology, the IA definition can be expressed as follows: IA ≡ Current and hasPhase A.
- Similarly, the Energy Ontology refers to a current of phase A as CA, then its definition (by reuse of the cross-domain concepts) becomes CA ≡ Current and hasPhase A.

A multi-layer approach enables great flexibility when querying since the ontologies are interconnected and queries can exploit both high level and specific concepts to explore a domain. Applications operating at a high-level of abstraction can use more general concepts to retrieve information such as (Current and hasPhase A), while applications operating at a more granular level can rely on the vertical ontologies to formulate queries and extract specific information such as CA and IA.

Taking a further example, oneM2M provides a set of rules to map the conceptual model of the oneM2M Base Ontology to the underlying oneM2M resource structure. Then for systems using other ontologies for which a mapping to the oneM2M Base Ontology can be defined, this provides a mechanism for those (other) ontologies to be instantiated within a oneM2M system as resources

[5] http://project-haystack.org

with associated semantic annotation. This enables different IoT systems to interwork with each other via a common upper ontology (Base Ontology) and a common architecture.

Ontologies and Semantically Augmented Things

Designing ontologies is the first step towards the interoperability vision; the second step consists of enabling the sensors, devices and systems to express their contextual information and data by applying the ontologies.

The connected world is a diverse ecosystem comprising elements which range from small, constrained devices and sensors, to larger more complex modules and machinery. This diversity is reflected in the processing, storage and communication capabilities provided by the respective elements and it follows that the degree to which ontologies and semantic capabilities can be embedded will also vary.

These considerations impact constrained elements for which embedding semantics is not an option. Sparsely resourced sensors, for example, will often provide little in terms of processing and may only support very lightweight, near binary format communications. In these cases, metadata can be (externally) attached to the sensor's data in a process referred to as semantic annotation [19] [20] . Semantic annotation is usually performed by the agent receiving the sensor's data, for example a gateway, a system, or a cloud agent [20] [21] .

We distinguish between two semantic annotation mechanisms; automatic tagging and commissioning [19] .

Automatic Tagging can be handled by a software agent running on a gateway, on a system or in the cloud. The agent decodes the sensor data stream (for example) and then augments the data using an appropriate semantic representation - see Figure 8. Other approaches, such as in [22] , suggest a heuristic based inference to harmonize the tags based on previously existing unstructured data.

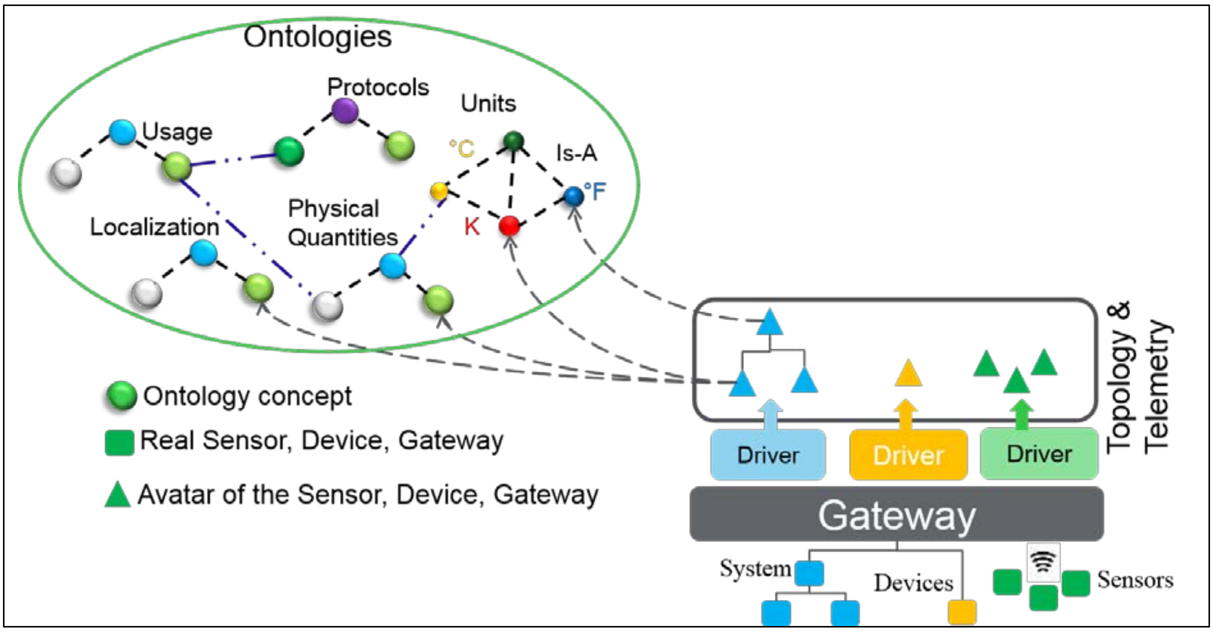

Figure 8 Semantic Annotation

Commissioning is usually handled through a user interface during the installation phase of a gateway or a system. For example, the use and location of a given sensor are only known during the commissioning phase and at that time the installer uses a commissioning tool to set the data from the ontologies. Commissioning tools should evolve to take into account such tagging.

Depending on the resources of the gateway or system, such annotation can remain as tags which can be processed by query engines to answer specific queries, as in [19] . Other approaches can rely on such tags to generate a complete ontology, as in [20] [21] .

Ontologies and Context

IoT technology is driving new opportunities for context-aware systems and applications. These classes of sentient system and application are able to adapt their behavior to the current context without explicit intervention.

Context awareness often means that a system combines physical awareness (time, location, sound, movement, touch, temperature) with application awareness (tasks, goals, processes, compliance, compatibility, approval, user disposition) to modify its own behavior.

Metadata alone has proven insufficient to address interoperability; in some institutions it is not part of software engineering curricula. Due to their formal expressiveness and the possibilities for applying ontology reasoning techniques, various existing and emerging context-aware frameworks use ontologies in their implementation [23] . Context sensitive machine learning techniques are also finding a role in deriving interoperability contexts [24] .

Ontologies and Scalability

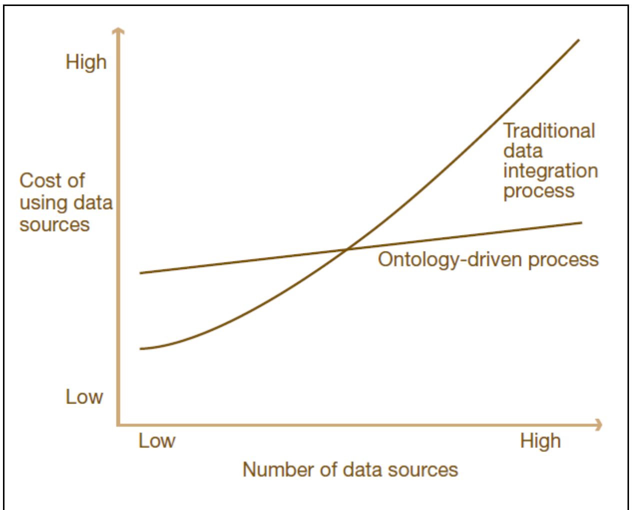

Source: PricewaterhouseCoopers, 2009 [25]

Figure 9 Scalability of ontology-based integration

A growing number of devices and applications are delivering data-streams and events on a continuous basis. This growth in data volume and velocity is accompanied by a growth in variety, driven by the heterogeneity of device and data formats.

Data integration programs based on traditional approaches such as relational databases are efficient in small, static and closed environments. In fact, with a low number of data sources, the cost of using and maintaining data remains low compared to an ontology-driven approach where a larger initial amount of effort is required therefore involving high costs.

The benefit of semantic standards stands out, however, when it comes to large, dynamic and open systems with critical requirements in terms of scalability and interoperability. Semantic models enable integrating a huge number of heterogeneous and mobile sources in short period of time with reasonable costs compared to traditional approaches. Semantic integration is performed once at the beginning, paving the way for advanced querying and reasoning, and enabling data to be integrated in a collaborative, standard and reusable way.

Current and Emerging Practices

Technologies and Strategies for Linked Data

The key to overcoming the fragmentation of the IoT and catalyzing exponential growth in services will be enabling end-to-end interoperability across different platforms. This requires open standards for metadata that define the data and interaction models exposed to applications, the protocols involved, and the communication patterns that can be used. In other words, this requires standards for a web of things.

What are the technologies and strategies for handling such metadata?

The Resource Description Framework (RDF) [26] provides for globally unique identifiers for metadata. These identifiers in many cases serve as links to further information for a web of *linked data*. RDF allows data and metadata to be described in terms of *triples*, i.e. named relationships that connect a subject to an object. There are multiple serialization formats for RDF, e.g. RDF-XML [27] , Turtle [28] and N3 [29] , comma separated values [30] , and JSON-LD [31] . Further techniques address how to include metadata within web pages.

Semantic models can be expressed with RDF Schema (RDF-S) [32] or the Web Ontology Language (OWL) [33] . SPARQL [34] is a query language for accessing and updating RDF triples. The Linked Data Platform (LDP) [35] defines how to use HTTP for read-write linked data on the web. DCAT [36] is an RDF vocabulary designed to facilitate interoperability between data catalogs published on the Web. Linked Open Vocabularies (LOV) community, introduced in the "Sharing Metadata" section, maintains descriptions of RDF-S vocabularies and OWL ontologies used for datasets in the Linked Data Cloud, see [37] .

Semantics in Support of Cross-Domain IoT

Semantic technologies provide a common means to describe domain knowledge whilst enabling heterogeneity and multimodality through interoperable data formats and various semantic models [38] .

Beyond the representational aspects, semantic computing supports reasoning on raw sensor data enabling the derivation of higher level abstractions; such abstractions form the basis of domain- and cross-domain knowledge. The Machine-to-Machine Measure (M3) Framework discusses these concepts [41] and provides an implementation which explores the creation of cross-domain applications.

Figure 10, which is discussed in [42] , outlines how M3 applies a semantic approach to enable cross-domain reasoning. The figure shows two sensors in different domains; sensor A is in the health domain and sensor B is in the weather domain.

Designing Semantic Models

Semantic data model design can be mainly split into two phases:

1. *Specification*, i.e. the conceptual/logical/abstract definition of the model. This is mainly mapped out in the form of objects, materialized as classes that can be linked together (relationships) and that are described by attributes
2. *Formalization*, i.e. its physical model in the form of semantic metadata or ontology.

Conceptual models are generally structured through the use of Entity Relationship (ER) or Unified Modeling Language (UML) diagrams. In line with practices commonly followed in Semantic Web [43] and RDF [44] development, Java conventions [45] are generally observed. For example:

- Adopt naming conventions which impose minimum changes
- Replace spaces in strings with underscores
- Use lower-case for metadata and ontologies namespaces
- Use camel-case for class and object names
- Use mixed-case for property names

The formalization of a semantic data model is achieved through the use of description languages.

Lightweight description languages such as JSON-LD (JavaScript Object Notation for Linked Data) [46] are generally preferred when dealing with low-power, low-energy, constrained embedded devices such as sensors and actuators. Less constrained environments commonly use XML based description languages such as OWL [47] . Best practices are available through the Semantic Web and Linked Data initiatives [48] [49] .

As discussed earlier, an ability to derive cross-domain value in an IoT context requires an ability to leverage multiple ontologies. *Leveraging* implies addressing functions such as ontology publication, discovery, reuse and mapping. Best practices for ontology engineering are available, for example [50] , but there are no *de facto* or standard references.

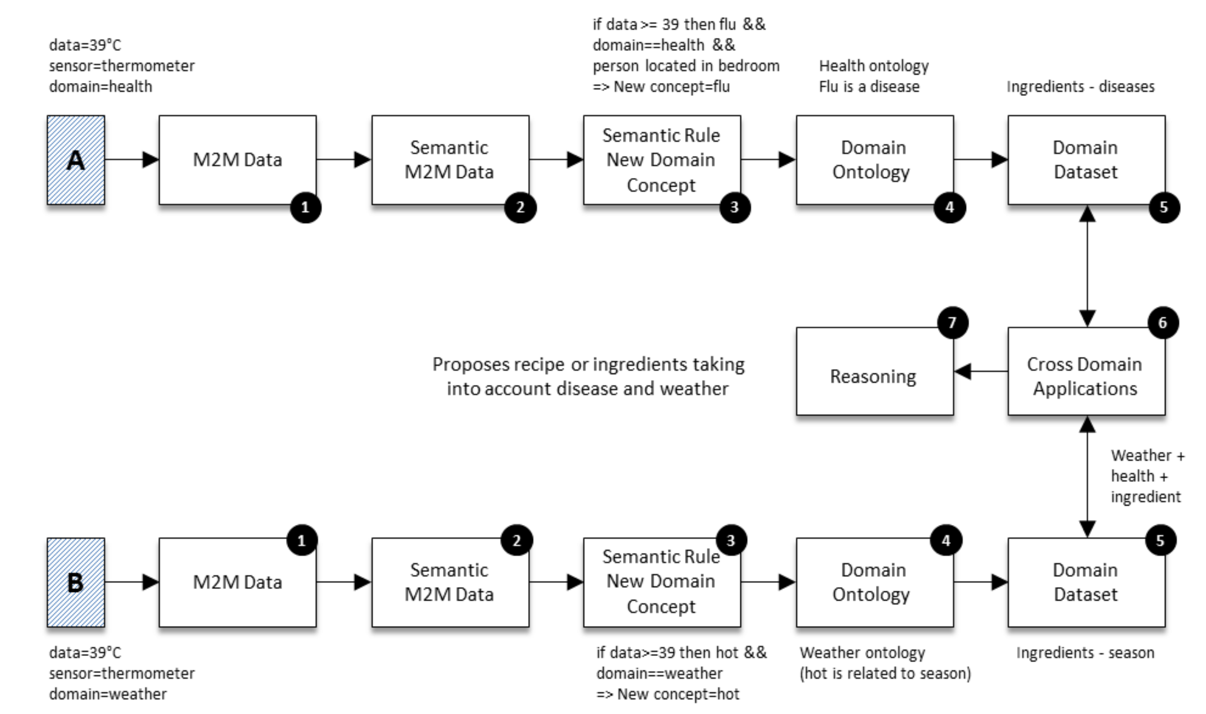

Step 1:	The raw measurements generated by the sensors are transformed into metadata which contain additional attributes: ▪ Unit of Measurement ▪ Timestamp ▪ Software Version ▪ Name ▪ Type ▪ Domain of Operation
Step 2:	The framework encodes the metadata using Sensor Markup Language [39] before converting into RDF [40] to enable semantic reasoning.
Step 3:	Semantic reasoning drives higher level abstractions as new domain concepts. In the health domain the reasoning leads to the concept of “flu”; in the weather domain to the concept of “hot”.
Step 4:	The respective domain ontologies are used to classify these new concepts; “flu” as a disease and “hot” as a seasonal condition.
Step 5-7:	The concepts, rules and datasets of the two domains are combined and cross-domain semantic reasoning takes place. In this example, the cross-domain reasoning produces suggestions for recipes appropriate for a given state of health and the prevailing weather conditions. The suggestions can be acted upon both by end-users and intelligent machines.

Figure 10 Enabling cross-domain scenarios using the M3 framework

The Emergence of API First and Microservices

Many of the major cloud service and software providers are giving attention to *API First* [51] application development approaches and *serverless* computing architectures [52] . Built on both in-house technologies and technologies from innovative startups these strategies target much of the heavy lifting in terms of hardware and software infrastructure (that includes provisioning for availability, throughput management, runtime analytics, etc.) enabling developers to focus the majority of their attention on the specific functionality of their respective products.

A further trend is the growing application of the *microservices* architectural style "an approach to developing a single application as a suite of small services, each running in its own process and communicating with lightweight mechanisms, often an HTTP resource API" [53] .

API First, serverless computing and microservices are being driven by the growing penetration of smart devices (in homes, cars, buildings, etc.), the ubiquity of smart phones and tablets as powerful mobile computing platforms, and the pressure to deliver the methods, tools and infrastructure to enable rapid realization of sophisticated, secure and scalable cloud-based IoT applications. The framework-oriented *Ontology as a Service* [54] or microservices conforming to the SSN-based IoT-Lite [55] are illustrative of these approaches.

Given the inherent cross-domain nature of IoT, it's likely that these forces will combine to accelerate the move towards semantic interoperability.

Conclusion and Closing Position

Although both the deployment of semantic technologies and the availability of skills are at fairly embryonic stages, there is a growing understanding that shared approaches to semantic interoperability is one of the keys to unlocking the value potential of the IoT. Indeed, this paper, which results from the efforts of a broad range of individuals responding to calls within the AIOTI WG3, oneM2M, IEEE P2413 and W3C communities, underlines that fact.

Realizing semantic interoperability at scale will require collaboration and coordination across standards organizations, consortia, alliances, and open source projects. The need for a shared roadmap and commitment to work together seems self-evident.

An initial focus may be created around lightweight models of semantics sufficient to manage interoperability of common domain independent and domain specific terms. There is increasing interest in providing agile processes for standardizing such terms and defining accessible, usable schemes for discovery and reuse. Such schemes will need to address different stages of lifecycle and maturity, e.g. from experimental, to commercial implementations, to (eventual) deployment on a global scale. An interesting precedent is provided by schema.org which defines a widely used lightweight RDF compatible vocabulary for websites to describe themselves to search engines.

Could this be generalized to descriptions of IoT services?

We challenge the community to initiate a process of alignment, consolidation and focus around semantic interoperability by:

- Establishing focused collaborations
- Creating a shared roadmap for addressing semantic interoperability concerns
- Adopting the roadmap to identify priorities, inform programs and rationalize deliverables

We look to standards organizations, consortia, alliances, and open source projects to endorse this approach and to proactively move the agenda forward.

References

[1] Van der Veer, H., Wiles, A. *Achieving Technical Interoperability - the ETSI Approach*. 3rd Ed. 2008. ETSI.
[Online] Available at:
http://www.etsi.org/images/files/ETSIWhitePapers/IOP%20whitepaper%20Edition%203%20final.pdf
[Accessed 22 Aug 2016]

[2] W3C Semantic Integration & Interoperability Using RDF and OWL
[Online] Available at:
https://www.w3.org/2001/sw/BestPractices/OEP/SemInt/
[Accessed 22 Aug 2016]

[3] IoT Analytics GmbH. *Disambiguation IoT vs IoE vs M2M vs Others*. 2015.
[Online] Available at:
http://iot-analytics.com/internet-of-things-definition
[Accessed 22 Aug 2016]

[4] Beecham Research Ltd. *M2M Sector Map*. 2009.
[Online] Available at
http://www.beechamresearch.com/download.aspx?id=18
[Accessed 22 Aug 2016]

[5] McKinsey Global Institute. *The Internet of Things: Mapping the Value Beyond the Hype*, 2015.
[Online] Available at
https://www.mckinsey.de/sites/mck_files/files/unlocking_the_potential_of_the_internet_of_things_full_report.pdf
[Accessed 22 Aug 2016]

[6] oneM2M TR-0001-V1.8.1, *Use Cases Collection*.
[Online] Available at
http://www.onem2m.org/technical/latest-drafts
[Accessed 22 Aug 2016]

[7] oneM2M Partners. *Study of Abstraction and Semantics Enablement.* 2014. oneM2M Technical Report TR-0007-V1.0.0.
[Online] Available at
http://www.onem2m.org/images/files/deliverables/TR-0007-Study_on_Abstraction_and_Semantics_Enablement-V1_0_0.pdf
[Accessed 22 Aug 2016]

[8] European Commission. *Horizon 2020 Work Programme 2016-2017: Internet of Things Large Scale Pilots*. 28-October-2015.
[Online] Available at
https://ec.europa.eu/digital-single-market/en/news/horizon-2020-work-programme-2016-2017-internet-things-large-scale-pilots
[Accessed 22 Aug 2016]

[9] W3C Semantic Sensor Network Incubator Group. *Semantic Sensor Network Ontology*.
[Online] Available at
https://www.w3.org/2005/Incubator/ssn/ssnx/ssn
[Accessed 22 Aug 2016]

[10] oneM2M Partners. Base Ontology. 2016. oneM2M Technical Specification TS-0012-V0.10.0. [Online] Available at
http://www.onem2m.org/component/rsfiles/download-file/files?path=Release_2_Draft_TS%255CTS-0012-oneM2M-Base-Ontology-V0_10_0.doc&Itemid=238
[Accessed 22 Aug 2016]

[11] ETSI. *Smart Appliances; Reference Ontology and oneM2M Mapping*. 2015. ETSI TS 103 264 V1.1.1.
[Online] Available at
http://www.etsi.org/deliver/etsi_ts/103200_103299/103264/01.01.01_60/ts_103264v010101p.pdf
[Accessed 22 Aug 2016]

[12] A.Gyrard, C.Bonnet,K.Boudaoud. *Domain knowledge Interoperability to build the Semantic Web of Things W3C Workshop on the Web of Things*, 25-26 June 2014.
[Online] Available at
https://www.w3.org/2014/02/wot/papers/gyrard-1.pdf
[Accessed 22 Aug 2016]

[13] P. Brandt et al., *Semantic interoperability in sensor applications making sense of sensor data*, Computational Intelligence in Healthcare and e-health (CICARE), 2013 IEEE Symposium on, Singapore, 2013, pp. 34-41.

[14] Cuenca Grau, B., Horrocks, i., Kazakov, Y., Sattler, U. Modular reuse of ontologies: Theory and practice. Journal of Articial Intelligence Research vol. 31, pp. 273-318, 2008.
[Online] Available at
http://www.cs.ox.ac.uk/people/yevgeny.kazakov/publications/journ/CueHorKazSat08Modularity_JAIR.pdf
[Accessed 22 Aug 2016]

[15] d'Aquin, M., Schlicht, A., Stuckenschmidt, H., Sabou, M. "Ontology modularization for knowledge selection: Experiments and evaluations," In Proceedings of DEXA2007, pp. 874-883, Regensburg, Germany, September 2007

[16] Doran, P. Ontology Modularization: Principles and Practice. Doctoral thesis, University of Liverpool, 2009.
[Online] Available at
https://core.ac.uk/download/files/129/105104.pdf
[Accessed 22 Aug 2016]

[17] ETSI TS 103 378 (V1.1.1) (10-2015): "Smart Body Area Network (SmartBAN); Unified data representation formats, semantic and open data model ".
[Online] Available at
http://www.etsi.org/deliver/etsi_ts/103300_103399/103378/01.01.01_60/ts_103378v010101p.pdf
[Accessed 22 Aug 2016]

[18] Nachabe, L., Girod-Genet, M., El Hassan, B. Unified data model for Wireless Sensor Network. IEEE Sensors Journal, vol. 15, n° 7, pp. 3657-3667, 2015.

[19] Charbel El Kaed and I. Khan and H. Hossayni and P. Nappey, SQenIoT: Semantic Query Engine For Industrial Internet-Of-Things Gateways, Submitted IEEE GLOBECOM 2016.

[20] C. Victor et al., An ontology design pattern for IoT device tagging systems, 5th International Conference on the Internet of Things, 2015.

[21] Charbel El Kaed, Brett Leida, Tony Gray, Building Management Insights Driven by a Multi-System Semantic Representation Approach, Submitted IEEE IoT World Forum 2016.

[22] J. Gao et al., A data-driven meta-data inference framework for building automation systems, in Proceedings of the 2Nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments, BuildSys ’15. ACM, 2015, pp. 23–32.

[23] Baldauf, M., Dustdar, S. and Rosenberg, F. *A survey on context-aware systems*, 2007. Int. J. Ad Hoc and Ubiquitous Computing, Vol. 2, No. 4, pp.263–277

[24] Gonzalez, A.J., Brézillon, P. *Context in Computing: A Cross-Disciplinary Approach for Modeling the Real World*. 2014. Springer New York.

[25] Price Waterhouse Coopers. *Technology Forecast*. 2009.
[Online] Available at
http://franz.com/agraph/cresources/white_papers/PriceWaterhouse-report.pdf
[Accessed 22 Aug 2016]

[26] W3C Resource Description Framework (RDF).
[Online] Available at
https://www.w3.org/RDF/
[Accessed 22 Aug 2016]

[27] W3C RDF/XML Syntax Specification
[Online] Available at
http://www.w3.org/TR/REC-rdf-syntax/
[Accessed 22 Aug 2016]

[28] W3C RDF 1.1 Terse RDF Triple Language (Turtle). 2014
[Online] Available at
http://www.w3.org/TR/turtle/
[Accessed 22 Aug 2016]

[29] W3C Notation3 (N3): A readable RDF syntax
[Online] Available at
http://www.w3.org/TeamSubmission/n3/
[Accessed 22 Aug 2016]

[30] W3C Comma Separated Values and RDF
[Online] Available at
https://www.w3.org/2013/csvw/wiki/Main_Page
[Accessed 22 Aug 2016]

[31] W3C A JSON-based Serialization for Linked Data. 2014
[Online] Available at
http://www.w3.org/TR/json-ld/
[Accessed 22 Aug 2016]

[32] W3C RDF Schema 1.1 (RDF-S). 2014
[Online] Available at
http://www.w3.org/TR/rdf-schema/
[Accessed 22 Aug 2016]

[33] W3C Web Ontology Language (OWL)
[Online] Available at
https://www.w3.org/OWL/
[Accessed 22 Aug 2016]

[34] W3C SPARQL 1.1 Query Language. 2013
[Online] Available at
http://www.w3.org/TR/sparql11-query/
[Accessed 22 Aug 2016]

[35] W3C Linked Data Platform 1.0. 2015
[Online] Available at
http://www.w3.org/TR/ldp/
[Accessed 22 Aug 2016]

[36] W3C Data Catalog Vocabulary (DCAT). 2014
[Online] Available at
http://www.w3.org/TR/vocab-dcat/
[Accessed 22 Aug 2016]

[37] Linked Open Vocabularies (LOV)
[Online] Available at
http://lov.okfn.org/dataset/lov/
[Accessed 22 Aug 2016]

[38] Datta, S. K., Gyrard, A., Bonnet, C., Boudaoud, K. *oneM2M Architecture Based User Centric IoT Application Development*. Future Internet of Things and Cloud (FiCloud), 2015 3rd International Conference on, Rome, 2015, pp. 100-107.

[39] Open Geospatial Consortium Inc.OpenGIS® Sensor Model Language (SensorML) Implementation Specification. 2007. OGC 0700 V1.0.0.
[Online] Available at
https://portal.opengeospatial.org/files/?artifact_id=55939
[Accessed 22 Aug 2016]

[40] W3C, *RDF 1.1 Concepts and Abstract Syntax,* 2015.
[Online] Available at
https://www.w3.org/TR/rdf11-concepts/
[Accessed 22 Aug 2016]

[41] Gyrard, A., Datta, S. K., Bonnet, C., Boudaoud, K. *Cross-Domain Internet of Things Application Development: M3 Framework and Evaluation*. Future Internet of Things and Cloud (FiCloud), 2015 3rd International Conference on, Rome, 2015, pp. 9-16.

[42] Gyrard, A., Bonnet, C., Boudaoud, K. *Enrich Machine-to-Machine Data with Semantic Web Technologies for Cross-Domain Applications*. IEEE World Forum on Internet of Things (WFIoT), pp. 559-564, 6-8 March 2014.

[43] W3C WoT Current Practices.
[Online] Available at
http://w3c.github.io/wot/current-practices/wot-practices.html
[Accessed 22 Aug 2016]

[44] Martin, P. *Propositions of Conventions for RDF,* 2000.
[Online] Available at
http://www-sop.inria.fr/acacia/personnel/phmartin/RDF/conventions.html
[Accessed 22 Aug 2016]

[45] Oracle. Naming conventions, 1999.
[Online] Available at
http://www.oracle.com/technetwork/java/codeconventions-135099.html
[Accessed 22 Aug 2016]

[46] See [31]

[47] See [33]

[48] Berrueta, D. and Phipps, J. *Best Practice Recipes for Publishing RDF Vocabularies - W3C Working Group Note*, 2008.
[Online] Available at
http://www.w3.org/TR/swbp-vocab-pub/
[Accessed 22 Aug 2016]

[49] Hyland, B., Atemezing, G., Villazón-Terrazas, B. *Best Practices for Publishing Linked Data, 5: The Role of "Good URIs" for Linked Data.* , 2014.
[Online] Available at
http://www.w3.org/TR/ld-bp/#HTTP-URISBest Practices for Publishing Linked Data
[Accessed 22 Aug 2016]

[50] Bergman, M. *A Reference Guide to Ontology Best Practices*. 2010.
[Online] Available at
http://www.mkbergman.com/911/a-reference-guide-to-ontology-best-practices/
[Accessed 22 Aug 2016]

[51] Kapoor, C. *The Future Of Coding Is Here, And It Threatens To Wipe Out Everything In Its Path*. 2015. Techcrunch. New York, NY: AOL Inc. Retrieved from citeulike-article-id:14074444
[Online] Available at
https://techcrunch.com/2015/09/27/the-future-of-coding-is-here-and-threatens-to-wipe-out-everything-in-its-path/
[Accessed 22 Aug 2016]

[52] Knorr, E. InfoWorld. *What serverless computing really means*
[Online] Available at
http://www.infoworld.com/article/3093508/cloud-computing/what-serverless-computing-really-means.html
[Accessed 22 Aug 2016]

[53] Fowler, M. *Microservices: A Definition of This New Architectural Term.*
[Online] Available at
http://martinfowler.com/articles/microservices.html
[Accessed 22 Aug 2016]

[54] Flahive A, Taniar D, Rahayu W. *Ontology as a Service (OaaS): extending sub-ontologies on the cloud*. Concurr Comput Pract Exp. 2015;27(8):2028-2040. doi:10.1002/cpe.3404.

[55] Bermudez-Edo M, Elsaleh T, Barnaghi P, Taylor K. *IoT-Lite Ontology, A Member Submission*. Cambridge, MA: W3C; 2015.
[Online] Available at
https://www.w3.org/Submission/2015/SUBM-iot-lite-20151126/
[Accessed 22 Aug 2016]