

The SSN Ontology of the W3C Semantic Sensor Network Incubator Group

Michael Compton^a, Payam Barnaghi^b, Luis Bermudez^c, Raúl García-Castro^d, Oscar Corcho^d, Simon Cox^e, John Graybeal, Manfred Hauswirth^f, Cory Henson^g, Arthur Herzog^h, Vincent Huangⁱ, Krzysztof Janowicz^{j,k}, W. David Kelsey^l, Danh Le Phuoc^f, Laurent Lefort^a, Myriam Leggieri^f, Holger Neuhaus, Andriy Nikolov^l, Kevin Page^m, Alexandre Passant^f, Amit Sheth^g, Kerry Taylor^a

^aCSIRO ICT Centre, Australia

^bUniversity of Surrey, UK

^cOpen Geospatial Consortium, USA

^dUniversidad Politécnica de Madrid, Spain

^eCSIRO Earth Science and Resource Engineering, Australia

^fDERI at the National University of Ireland, Galway, Ireland

^gWright State University, USA

^hFraunhofer Gesellschaft, Germany

ⁱEricsson

^jUniversity of California, Santa Barbara, USA

^k52° North Initiative for Geospatial Open Source Software, Münster, Germany

^lThe Open University, UK

^mUniversity of Southampton, UK

Abstract

The W3C Semantic Sensor Network Incubator group (the SSN-XG) produced an OWL 2 ontology to describe sensors and observations — the SSN ontology, available at <http://purl.oclc.org/NET/ssnx/ssn>. The SSN ontology can describe sensors in terms of capabilities, measurement processes, observations and deployments. This article describes the SSN ontology. It further gives an example and describes the use of the ontology in recent research projects.

Keywords: Semantic Sensor Web, Linked Sensor Data, Ontologies, Semantic Interoperability, Sensor Web Enablement

1. Introduction

Observations, and the sensors that obtain them, are at the core of empirical science. Sensors are used in applications ranging from meteorology to medical care to environmental monitoring to security and surveillance. The use of sensing devices and networked sensing devices is increasing. This increase is accompanied by an increasing volume of data, as well as increasing heterogeneity of devices, data formats, and measurement procedures.

Therefore, as the prevalence of sensing devices and systems grows, ways to manage the sensors and accompanying volume of generated data become important. The Sensor Web Enablement (SWE) [1] initiative of the Open Geospatial Consortium (OGC) defined data encodings and Web services to store and access sensor-related data. These standards, for example, SensorML [2] and O&M [3, 4], provide syntactic interoperability [5]. An additional layer is required to address semantic compatibility [6].

Semantic Web technologies have been proposed as a means to enable interoperability for sensors and sensing systems. Semantic Web technologies could be used in isolation or in augmenting SWE standards in the form of the Semantic Sensor Web [5].

Semantic technologies can assist in managing, querying, and combining sensors and observation data. Thus allowing users to operate at abstraction levels above the technical details of format and integration, instead working with domain concepts and restrictions on quality. Machine-interpretable semantics allows autonomous or semi-autonomous agents to assist in collecting, processing, reasoning about, and acting on sensors and their observations. Linked Sensor Data [7, 8] may serve as a means to interlink sensor data with external sources on the Web.

Shared semantic definitions help not only with data integration from multiple sources, but can also assist in integrating new data into historical, temporal and spatial contexts. Definitions of sensors and their capabilities are also useful for provenance and quality reasoning. For example, if the accuracy of a sensor depends on phenomena other than that which it measures, then a specification of this can be used as a guide to search for spatially and temporally related measurements of the phenomena on which the accuracy depends, allowing the calculation of quality metrics.

The W3C Semantic Sensor Network Incubator group (SSN-XG) defined an OWL 2 [9] ontology to describe the capabilities and properties of sensors, the act of sensing and the resulting observations. This article describes the ontology, its development, an example and current uses.

Email address: Michael.Compton@csiro.au (Michael Compton)

Paper Contributions and Organisation

This paper makes two contributions. The bulk of the paper relates to the first and central contribution: a description of the SSN ontology (§2 – §7). An outline of the SSN-XG (§2) and development of the ontology (§2) is presented first, followed by the general structure of the ontology (§3). The detailed description of the ontology begins with a central pattern (§4), followed by sensor (§5), observation (§6) and system (§7) perspectives.

The second contribution is an example of the use of the ontology (§8.1) and a discussion of projects and applications in which it has been used (§8.2). This is followed by concluding remarks (§9).

2. The W3C Semantic Sensor Network Incubator Group

W3C incubator groups are one-year (plus possible extensions) exploratory activities on emerging Web-related concepts, guidelines or activities. They can lead to further W3C activities, member submissions or recommendations. The SSN-XG was initiated by the CSIRO, Wright State University, and the OGC as a forum for the development of an OWL ontology for sensors and to further investigate annotation of, and links to, existing standards.

The SSN-XG ran from March 2009 to September 2010. Some 41 people, from 16 organisations, joined the group, with 20 members attending more than 10 meetings. The activities of the group were recorded on the group’s wiki,¹ where the group’s final report² can be found. The final report includes sections on use cases, the group’s review of existing sensor and observation related ontologies, the SSN ontology, mappings of terms from the ontology to other standards and vocabularies, and material on the group’s other main deliverable on semantic annotation of OGC Sensor Web Enablement standards.

Development of the SSN Ontology

The group began by reviewing existing ontologies and standards, and developing use cases. The use cases were focused into four categories — entitled: data discovery and linking, device discovery and selection, provenance and diagnosis, and device operation, tasking and programming — with a prototypical example in each.

The device discovery and selection use case, for example, requires the ontology to represent sensor types, models, methods of operation and common metrological definitions like accuracy, precision, measurement range, and the like, thus allowing sensor capabilities to be defined relative to prevailing conditions. Such definitions would enable a sensor network designer to search a database for sensors matching required parameters of operation and accuracy, perhaps cross checking against recorded climatic conditions or running a simulator to analyse expected performance of such a network.

The data discovery and linking use case focused on finding and linking data, given qualitative, spatial or temporal bounds. It requires the sensor specifications from the device discovery and selection case, observation data and linkage to other data sources. While the provenance and diagnosis use case requires context information from sensor and observation data, deployment information, custodian descriptions, maintenance schedules and data linkage to derive trust levels or to further analyse previous measurements. The device operation, tasking and programming use case requires sufficient information to reprogram a device or understand the consequences, in terms of, say, energy usage or network cost, of a reprogramming.

In documenting the expectations of group members, such use cases also serve to expose the tension between completeness and focus in ontology modelling: the choice between an ontology that can model programming concepts, observed phenomena and maintenance schedules and one that is sensor only. Indeed, the review (see also Compton et al. [10]) showed that while there were a number of existing artefacts (twelve ontologies were reviewed) none covered the requirements of the use cases or satisfied a design goal of limiting to sensor specific concepts and relations.

Limiting the ontology to sensor only concepts should serve to increase modularity and reusability and became the inclusion criteria for the group. Thus, the ontology should enable the sensor aspects of the use cases, without needing to fulfil all the modelling requirements.

Discussions revealed different interpretations of concepts; even fundamental concepts like sensor — a single sensing device, or any sensing system; a single stimulus to observation sensor, or one that allows multiple calculations and combination of stimuli.

Thus, after much discussion, and a few false starts, the group consensus was to build an ontology to describe sensors, that was: as far as possible sensor specific; compatible with OGC standards, without being constrained by them; and that generally chose the broadest definition for concepts, so that subconcepts could be later defined for more specific interpretations. For example, the result for sensor was any entity capable of following some method to sense, allowing sensing devices as a subconcept.

The SSN ontology was developed by group consensus over a period of one year. First, the core concepts and relations were developed (sensors, features and properties, observations, and systems). Then, measuring capabilities, operating and survival restrictions, and deployments were added in turn. Finally, the alignment to DOLCE-UltraLite³ (DUL) and the realisation of the core Stimulus-Sensor-Observation ontology design pattern [11] were added.

The alignment to a foundational ontology makes ontological commitments explicit, further explains concepts and relations, and restricts possible interpretations towards their intended meaning [12]. DUL was chosen as the upper ontology

¹http://www.w3.org/2005/Incubator/ssn/wiki/Main_Page (last accessed 11th May 2012)

²<http://www.w3.org/2005/Incubator/ssn/XGR-ssn/> (last accessed 11th May 2012)

³<http://www.loa-cnr.it/ontologies/DUL.owl> (last accessed 11th May 2012)

because it is more lightweight than other options, whilst having an ontological framework and basis, for example with qualities, regions and object categories, that was consistent with the group’s modelling.

3. The SSN Ontology

The SSN ontology, available at

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

is organised, conceptually but not physically, into ten modules, Figure 1. The full ontology consists of 41 concepts and 39 object properties, directly inheriting from 11 DUL concepts and 14 DUL object properties.

The ontology can describe sensors, the accuracy and capabilities of such sensors, observations and methods used for sensing. Also concepts for operating and survival ranges are included, as these are often part of a given specification for a sensor, along with its performance within those ranges. Finally, a structure for field deployments is included to describe deployment lifetime and sensing purpose of the deployed macro instrument.

Related, but not sensor specific, material such as units of measurement, locations, hierarchies of sensor types, and feature and property hierarchies were left aside. Where appropriate, concepts were included to allow linking to such external ontologies. For example, an observation is of a particular property of a feature, where observations are fully described by the ontology, while feature and property are left as place holder concepts. The intention is that in building an ontology based on the SSN ontology, knowledge engineers would include the SSN ontology, suitable units, location and feature ontologies, and link these via subclassing or equivalence relations. This combination can then be used to describe a hierarchy of sensors relevant to the particular application. The SSN-XG wiki pages contain a number of illustrative examples.

The SSN ontology is built around a central Ontology Design Pattern (ODP) [13] describing the relationships between sensors, stimulus, and observations, the Stimulus-Sensor-Observation (SSO) pattern [11]. The ontology can be seen from four main perspectives:

- A sensor perspective, with a focus on what senses, how it senses, and what is sensed;
- An observation perspective, with a focus on observation data and related metadata;
- A system perspective, with a focus on systems of sensors and deployments; and,
- A feature and property perspective, focusing on what senses a particular property or what observations have been made about a property.

The ontology takes a liberally inclusive view of what a sensor is: *anything that observes*; and allows such sensors to be described at any level of detail, for example, allowing sensors

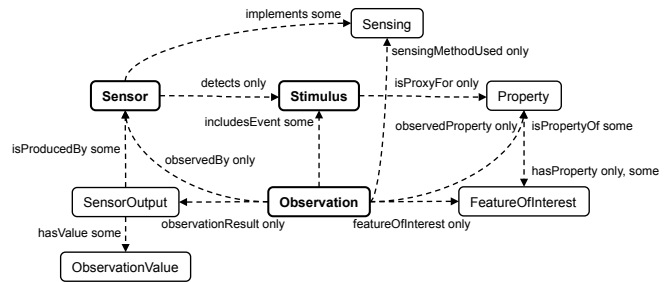


Figure 2: The Stimulus-Sensor-Observation Pattern (central concepts in bold).

to be seen simply as objects that play a role of sensing, as well as allowing sensors to be described in terms of their components and method of operation. Humans and also simulations can be modeled as sensors.

Concepts and properties in the ontology are commented with `rdfs:comment`, `rdfs:isDefinedBy`, `rdfs:label`, `rdfs:seeAlso` [14], and `dc:source`.⁴ The `rdfs:seeAlso` comments link to further explanation on the group’s wiki, while the `dc:source` comments use the SKOS [15] terms `skos:exactMatch` and `skos:closeMatch` to relate concepts and properties to SensorML, O&M and the International Vocabulary of Metrology [16, also known as ISO/IEC Guide 99:2007]. Scripts were developed that derive navigable documentation from the ontology.⁵

The following sections present the SSO pattern and the sensor, observation and system perspectives (the feature and property perspective being covered by the relations introduced in the other sections. Namespaces for the SSN and DUL ontologies are written prefixing concepts and properties as `ssn:` and `dul:`, respectively.

4. The Stimulus-Sensor-Observation Pattern

Central to the ontology is the Stimulus-Sensor-Observation ontology design pattern [11], Figure 2. The pattern links sensors, what they sense, and the resulting observations, encompassing three of the four perspectives — the missing system perspective is more about system organisation and deployments than sensing, but clearly links to the pattern. The SSO has been developed as a minimal, common ground for heavy-weight ontologies for the Semantic Sensor Web, as well as to explicitly address the need for light-weight semantics for Linked Data.

4.1. Stimuli

Stimuli are changes or states (`dul:Event`) in an environment that a sensor can detect and use to measure a property. A stimulus (`ssn:Stimulus`) is thus a proxy (`ssn:isProxyfor`) for an observable property (`ssn:Property`), or a number of observable properties. For example, changes in electrical resistance as

⁴Dublin core definitions are available at <http://dublincore.org> (last accessed 11th May 2012), where `dc:source` is defined as ‘A related resource from which the described resource is derived.’

⁵<http://www.w3.org/2005/Incubator/ssn/wiki/SSN> (last accessed 11th May 2012)




Figure 1: The SSN ontology, key concepts and relations, split by conceptual modules. The concepts not depicted are largely properties for measurement capabilities, and survival or operating ranges: accuracy, precision, resolution and the like. Note the central importance of sensors, observations and properties, brought out by the SSO ontology design pattern.

a proxy for temperature in a thermistor, or current generated by spinning wind cups for wind speed. Properties themselves are observable characteristics of (`ssn:isPropertyOf`) real-world entities (`ssn:FeatureOfInterest`). In the DOLCE alignment we specify `ssn:FeatureOfInterest` \sqsubseteq `dul:Event` \sqcup `dul:Object`, rather than using `dul:Entity`, since one cannot sense properties of abstract entities, such as sets and regions. The notion of a stimulus marks the borders of empirical science [17, 18].

4.2. Sensors

In the SSO ontology, sensors (`ssn:Sensor`) are physical objects (`dul:PhysicalObject`) that observe, transforming incoming stimuli (`ssn:detects`) into another, often digital, representation (`ssn:SensorOutput`). Sensors may be hardware devices, sensing systems, scientific computational models, human run laboratory setups — anything that senses. A sensor follows (`ssn:implements`) a method (`ssn:Sensing` \sqsubseteq `dul:Method` \sqsubseteq `dul:Description`) describing how the sensor observes: this may be, for example, a description of the scientific method implemented by the sensor.

The sensing method, though, is distinct from a process or workflow description of how the sensor operates (not described by the SSN ontology, workflow and process descriptions being more widely applicable and expected to be imported from a suitable ontology). A method is an abstract description; there may be any number of ways to concretely realise one.

4.3. Observations

Observations (`ssn:Observation`) are the nexus of the SSO pattern. For a sensing event, an observation can link the act of sensing (`dul:includesEvent`, not in the pattern), the event that is the stimulus (`dul:includesEvent`), the sensor (`ssn:observedBy`), a method (`ssn:sensingMethodUsed`), a result (`ssn:observationResult`), an observed feature (`ssn:featureOfInterest`), and property (`ssn:observedProperty`), placing all in an interpretative context. While observations have been modelled in different ways in the literature, we defined them as social constructs (`ssn:Observation` \sqsubseteq `dul:Situation`). That is, observations are contexts for interpreting incoming stimuli and fixing parameters such as time and location. Researchers interested in the internal processes by which sensors translate stimuli into other representations may model observations as events.

5. Sensor Perspective

The SSO pattern describes a sensor in terms of its stimulus, sensing method, and the observations it makes. The complete sensor perspective, Figure 3, enriches this picture to include the capabilities of sensors. For any property observed (`ssn:observes`) by a sensor, the performance (accuracy, etc) of the sensor might be affected by prevailing environmental conditions, related or not to the property under observation. This is modelled as measurement capabilities.




Figure 3: Ontology view focusing on measurement capabilities of sensors, with highlighting focus on measurement capabilities. A sensor may be linked to any number of capability descriptions, each specifying, for example, how the sensor’s accuracy and resolution are affected by prevailing environmental conditions.

The Measuring Capability of Sensors

The fact that the accuracy of a sensor is affected by prevailing conditions is an observable property of the sensor. Indeed, sensing devices, are often described by a data sheet that lists properties observed of the sensor in various conditions. That is, accuracy, measurement range, precision, resolution, and the like are all properties that one might observe of a sensor (which means that the capabilities of a sensor can be specified using observations recorded in the SSN ontology).

The ontology models `ssn:Accuracy`, `ssn:DetectionLimit`, `ssn:Drift`, `ssn:Frequency`, `ssn:Latency`, `ssn:MeasurementRange`, `ssn:Precision`, `ssn:ResponseTime`, `ssn:Resolution`, `ssn:Sensitivity` and `ssn>Selectivity` as measurement properties (`ssn:MeasurementProperty` \sqsubseteq `ssn:Property`). A sensor may have (`ssn:hasMeasurementCapability`) a number of measurement capabilities (`ssn:MeasurementCapability` \sqsubseteq `ssn:Property`), describing the capability of the sensor in (`ssn:inCondition`) various conditions (`ssn:Condition`), which are in turn observable properties of the sensor’s environment.

A measurement capability instance collects together observed properties of a sensor in the conditions specified. A sensor may have links, through `ssn:hasMeasurementCapability`, to any number of `ssn:MeasurementCapability` instances: for example, multiple instances with the same property, but different conditions, specify the capability of the sensor in a range of conditions, while, for sensors that observe multiple properties, multiple instances with different properties, perhaps each with different conditions, specify the capability of the various sensing functions.

Note that the SSN ontology does not restrict the observations of a device to fall within the prescribed measurement ranges (or make assertions about a device that is used beyond its measurement or survival capacity). These would be application specific algorithms to detect outliers, calculate quality metrics, determine damage, or any number of other parameters. Such algorithms would use the specification to determine the parameters (potentially asserting new facts back into the ontology) and may be implemented as OWL restrictions, rules, other code that access the knowledge base and reasoner through an API, or a

combination of these. The purpose of the ontology is to record the specification and enable such algorithms, not to enforce one such algorithm on all applications.

6. Observation Perspective

Complementing the sensor perspective is the observation perspective, which completes the description of an observation introduced in the SSO pattern. Observations are contexts for interpreting incoming stimuli and, hence, place the observing event and stimulus in an interpreting context. The context includes observed feature (`ssn:featureOfInterest`), property (`ssn:observedProperty`), observing sensor (`ssn:observedBy`), result (`ssn:observationResult`), and method (`ssn:sensingMethodUsed`) from the SSO pattern. It can also, in complement to a sensor’s capabilities, record an adjudged quality of the observation (`ssn:qualityOfObservation`), a time the result became available (`ssn:observationResultTime`) and a time at which the sampling took place (`ssn:observationSamplingTime`) - time being an aspect the SSN ontology does not describe and is left for an imported time ontology.

The treatment of an observation as a social construct, interpreting events, participants and associated result, differs from O&M, in which observations are seen as the observing events themselves, but is in line with the pattern proposed by Blomqvist.⁶ The treatment here has the benefit that it separates a stimulus event from potential multiple interpretations of it and that it signifies the interpretative nature of observing. Despite the different ontological classifications of observation, the associated data remains the same as O&M.

A sensing method can both describe the principle underlying a sensor and describe how observations were made: that is, the principle underlying the observation, describing, for example, how a sensor was positioned and used. In some cases, this allows a modelling choice, where, for example, sensing devices used in a particular way could be best modelled as sensors used as per an observation method, whereas, a more intricate setup may be more appropriately modelled as a sensor than observation method.

The ontology is thus designed to support modelling complex cases. For instance, the same sensing device, such as a specific thermometer, can be used to measure soil as well as ground temperature. Measurements for these observed properties arrive at different observation values and refer to different features of interest and, hence, cannot be combined. It is the measurement *procedures* that distinguishes both observations. For instance, air temperature is typically measured 2m above ground with a sensor protected from direct solar radiation. Consequently, the SSN ontology can be used to foster semantic interoperability by preventing agents and Web services from combining both temperature measures. Probst and Lutz have shown that this is not possible on the syntactic level [19].

⁶ <http://ontologydesignpatterns.org/wiki/Submissions:Observation> (last accessed 11th May 2012)




Figure 4: Ontology view showing systems, deployments, platforms and operating and survival conditions.

7. System Perspective

The system perspective, Figure 4, is constructed around a system (`ssn:System`) concept representing parts of sensing infrastructure. A system has components (`ssn:hasSubSystem`) which are systems. Systems, of which devices and sensing devices are sub concepts (`ssn:SensingDevice` \sqsubseteq `ssn:Device` \sqsubseteq `ssn:System`), have operating and survival ranges (`ssn:hasOperatingRange` and `ssn:hasSurvivalRange`), may be mounted on platforms (`ssn:onPlatform`) and may be deployed (`ssn:hasDeployment`).

7.1. Operating and Survival Restrictions

Prevailing environmental conditions may affect the performance of a sensor, similarly, a system or device may have a defined operating environment, and environmental extremes may exceed the capacity of a system to survive and make further observations. The general structure for describing operating and survival ranges is the same as for sensors and measurement capabilities, indeed they are observable properties of systems. The operating range (`ssn:OperatingRange`), describing characteristic of the environmental and other conditions in which the system is intended to operate, includes features such as power ranges, power sources, standard configurations, and attachments. The survival range (`ssn:SurvivalRange`) describes environmental conditions to which a sensor can be exposed without causing lasting damage: i.e., the sensor continues to operate as per defined measurement capabilities. If, however, the survival range is exceeded, the sensor is considered damaged such that measurement capability specifications may no longer hold.

7.2. Deployment

A deployment (`ssn:Deployment` \sqsubseteq `dul:Process` \sqsubseteq `dul:Event`) is a process that encompasses all phases in the life-time of a deployed system: such as, installation, maintenance and decommissioning. A system is deployed on (`ssn:deployedOnPlatform`) a platform (`ssn:Platform` - a role an entity plays whilst a system is attached).

Locations of platforms, systems or sensors and temporal properties of deployments are areas where other ontologies are required to fill in the details. Broadly, location can be represented as either abstractions of real-world locations, or as absolute or relative locations. In the first case, relating (`dul:hasLocation`) a sensor to a place (`dul:PhysicalPlace`), as in the sensor/platform is on the eastern edge of the lake; indeed the relation between a sensor and a platform can be specified in this way (`ssn:onPlatform` \sqsubseteq `dul:hasLocation`). Absolute and relative locations, on the other hand, acknowledge that location is an observable aspect of an entity, which thus may have a property (`ssn:hasProperty`) stating location using, for example, absolute or relative latitude and longitude.

Temporal properties could be included by specifying a date for deployment processes (`dul:hasEventDate`) or by including a time ontology, perhaps treating time as observable and classifying time concepts into the DUL hierarchy (using `dul:TimeInterval`).

8. Examples and Uses of the SSN ontology

8.1. Examples

Examples on the group's wiki combine the SSN ontology with units of measurement, feature and quality, and domain ontologies. The examples are: two LOD examples, from the SENSEI⁷ project [7] and the Kno.e.sis Center⁸ (also an example of semantically annotating streaming sensor data); a sensor embedded products example from the SmartProducts project;⁹ an agriculture and meteorology example; and a specification drawn from a commercial sensor data sheet.

The ontology of the data sheet example,¹⁰ Figure 5, imports the SSN ontology and the QU ontologies¹¹ of quantities and units. The example uses properties (called quantity kinds in QU), units of measure, and values from the QU ontologies. QU quantity kinds are specified as subsumed by `ssn:Property` and the QU units are used in specifying concrete values.

It is a good case study because the device measures two properties (wind speed and direction) and the accuracy of the wind speed measurements is dependent on prevailing conditions. That the accuracies are expressed as relative and absolute errors further shows the detail that can be obtained.

⁷<http://www.sensei-project.eu/> (last accessed 11th May 2012)

⁸<http://knoesis.wright.edu/> (last accessed 11th May 2012)

⁹<http://www.smartproducts-project.eu/> (last accessed 11th May 2012)

¹⁰<http://purl.oclc.org/NET/ssnx/meteo/WM30> (last accessed 11th May 2012)

¹¹<http://purl.oclc.org/NET/ssnx/qu/qu> and <http://purl.oclc.org/NET/ssnx/qu/qu-rec20> (last accessed 11th May 2012)

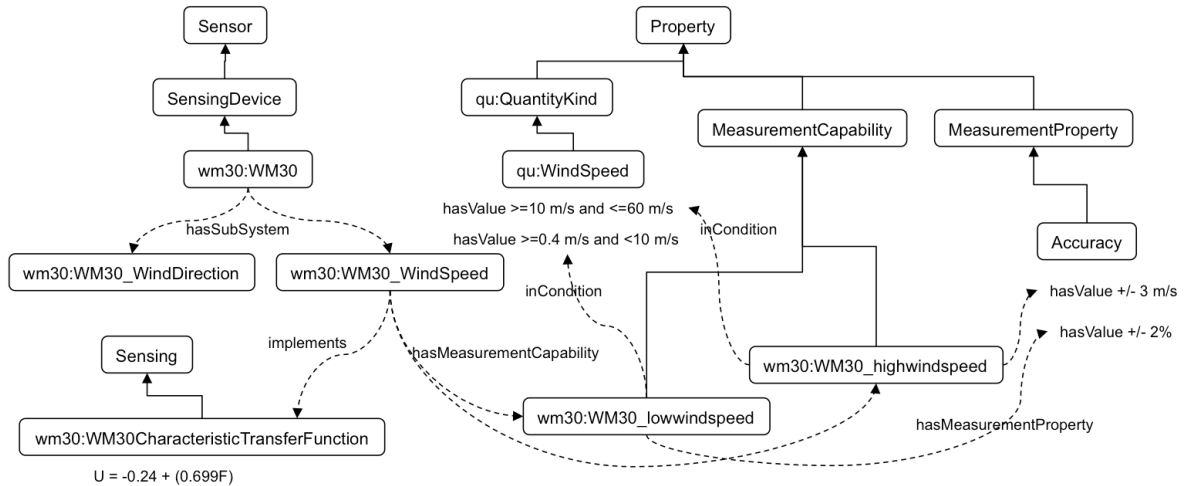


Figure 5: Simplified view of wind sensor example. The wind speed sensor has accuracy dependent on wind conditions. Not shown in the figure are the operating and survival ranges and the specification of the wind direction sensor, which comes with options for 355° and 360° measurement ranges. The sources of new concepts are shown with wm30: and qu: namespaces. Values incorporate DUL regions and QU units.

Further, the device data sheet also includes the characteristic transfer function of the device: $U = -0.24 + (0.699 \times F)$, where U is wind speed in m/s and F is frequency in Hz . This is the formula used by the device to relate the turning of the cups (the frequency) to windspeed: it is the measurement method. In the example, the function is documented as a string, but could equally be expressed in MathML [20], as a program fragment, or described using an ontology for processes and workflows, such as one based on OWL-S¹² or PML¹³ — such formal definitions can be used to construct sensors from specifications [21].

The example would help satisfy the device selection and provenance use cases, and could be extended with a suitable device ontology with concepts for size, weight, power usage and, if the device were reprogrammable, as in Taylor and Penkala [22], concepts for the tasking and reprogramming use case.

The example specifies properties of all such devices, not a particular device. Generally, an ontology extending the SSN ontology defines types and classes of devices as (TBox) concepts and instances (actual devices in particular places) as (ABox) individuals. This division, between TBox definitions of sensor classes and ABox definitions of sensor instances (similarly for observations), means ABox individuals need not repeat the general information for the sensor type, which is inferred to be present. Instead, only contextual information, such as location, deployment, etc., is required; an advantage over SensorML which has no such class/instance division.

8.2. Uses of the SSN Ontology

By people not involved in the SSN-XG, the ontology has been used as part of an architecture for the web of things [23], in sensing for manufacturing [24], for representing humans and

personal devices as sensors [25] and as part of a linked data infrastructure for SWE [26]

Along with those mentioned in Section 8.1, the ontology has also been used by SSN-XG participants in the following ways.

The SPITFIRE FP7 project¹⁴ combines the ontology with an event model,¹⁵ a context and cognition ontology and linked open data sources such as linked geodata¹⁶ and DBpedia.¹⁷ This gives a scalable data model for representing sensor information and inferred semantic descriptions of sensors, which is used, for example, with SPARQL in queries that involve both the sensor data and the contextual linked data. The linked data is also used to automatically annotate new sensors based on the context that can be discovered in the sensor data and the LOD cloud. For example, spatially or thematically related sensors can be discovered through the linked data with the integration of place names, coordinates, sensing types, etc. Linking sensed features, sensors and other Internet of Things objects through the LOD cloud gives a machine processable, but also human-friendly, semantic entity representation in line with the linked data paradigm (see also Leggieri et al. [27], Pfisterer et al. [28]).

The SSN ontology and the SSO pattern are also key for the Semantic Sensor Web and Linked Sensor Data work at 52°North.¹⁸ The ontology and SSO pattern are used in the implementation of a transparent and RESTful proxy for OGC's Sensor Observation Service (SOS). The proxy takes a URI and returns an RDF representation of requested observation data. The URIs are defined by a schema that provides both identity and filter encoding. The proxy can be installed in front of a SOS to serve Linked Sensor Data on-the-fly [29]. 52°North has also implemented a semantically-enabled *Sensor Plug & Play*

¹⁴<http://www.spitfire-project.eu/> (last accessed 11th May 2012)

¹⁵http://ontologydesignpatterns.org/wiki/Ontology:Event_Model_F (last accessed 11th May 2012)

¹⁶<http://linkedgeodata.org/> (last accessed 11th May 2012)

¹⁷<http://dbpedia.org/> (last accessed 11th May 2012)

¹⁸<http://52north.org/> (last accessed 11th May 2012)

¹²<http://www.w3.org/Submission/OWL-S/> (last accessed 11th May 2012)

¹³http://tw.rpi.edu/portal/Proof_Markup_Language (last accessed 11th May 2012)

framework [30]. The SSN ontology is combined with ontologies for observed properties, such as SWEET,¹⁹ and SWE documents are annotated to link to these ontologies. On installation of a new sensor, the framework can thus interpret the metadata advertised in the corresponding SensorML document and serve up observations as annotated O&M. Further, if a requirement for a particular type of stream has been requested as a template request from SWE services, the framework can resolve the semantics of the SWE descriptions, taking into account for example property or sensor hierarchies, and uses SWRL rules to match and convert compatible units. This allows on-the-fly integration of sensors with the Sensor Web, minimising administration efforts. Thus the manual interaction and mediation required to register new sensors can be reduced to a minimum and the probability of semantic mismatching can be reduced.

The ontology is a fundamental ontology in the SensorGrid4Env project,²⁰ which aims to build large semantic sensor network applications for environmental management [31, 32]. The project deployed APIs for developing web applications and mashups that use the semantics to simplify the configuration and deployment of services. The SSN ontology is linked with ontologies for coastal features, services and roles for emergency response,²¹ as well as SWEET, and United Kingdom mapping data.²² Using spatio temporal and geographic extensions to SPARQL, data can be queried from sources such as relational databases, XML stores, or triple stores. The system allows the discovery of data sources and services based on content and spatiotemporal coverage as well as integration of heterogeneous sensor and stored data sources.

The ontology was used in publishing data from the Spanish Meteorological Agency.²³ It is used in the Exalted project²⁴ for query management, event processing and communication. And used at CSIRO in research on sensor network installation, querying, programming and provenance.²⁵

9. Conclusion

The W3C Semantic Sensor Network Incubator Group developed an OWL ontology — the SSN ontology — for describing sensors. The SSN ontology can describe sensors, sensing, the measurement capabilities of sensors, the observations that result from sensing, and deployments in which sensors are used. The ontology covers large parts of the SensorML and O&M standards, omitting calibrations, process descriptions and data types, each not sensor specific.

The Stimulus-Sensor-Observation ontology design pattern can be used separately for the light-weight semantics preferred

by some Linked Data applications. It may be preferable to arrange the ontology in physically separate modules, with, for example, the SSO pattern, the bulk of the sensor concepts and deployments in separate ontologies (deployments ontology importing sensor concepts, which imports SSO, which imports DUL).

In allowing the TBox/ABox division discussed in Section 8.1, the SSN ontology allows class and instance definitions to be managed in separate ontologies, perhaps by separate authorities, an option not available for SensorML: for example, device manufacturers providing concept definitions for products and users specifying instances on deployment.

The SSN ontology is currently used in a number of research projects. Its development was informed largely by participants' use cases, existing OWL ontologies, OGC standards, and vocabularies such as the International Vocabulary of Metrology.

References

- [1] M. Botts, G. Percivall, C. Reed, J. Davidson, OGC Sensor Web Enablement: Overview And High Level Architecture, OpenGIS White Paper OGC 07-165, Open Geospatial Consortium Inc., 2007.
- [2] M. Botts, A. Robin, OpenGIS Sensor Model Language (SensorML) Implementation Specification, OpenGIS Implementation Specification OGC 07-000, The Open Geospatial Consortium, 2007.
- [3] S. Cox, Observations and Measurements – Part 1 – Observation schema, OpenGIS Implementation Standard OGC 07-022r1, Open Geospatial Consortium Inc., 2007.
- [4] S. Cox, Observations and Measurements – Part 2 - Sampling Features, OpenGIS Implementation Standard OGC 07-002r3, Open Geospatial Consortium Inc., 2007.
- [5] A. Sheth, C. Henson, S. Sahoo, Semantic Sensor Web, IEEE Internet Computing 12 (4) (2008) 78–83.
- [6] K. Janowicz, S. Schade, A. Bröring, C. Keßler, P. Maue, C. Stasch, Semantic Enablement for Spatial Data Infrastructures, Transactions in GIS 14 (2) (2010) 111–129.
- [7] P. Barnaghi, M. Presser, Publishing Linked Sensor Data, in: 3rd International Workshop on Semantic Sensor Networks, CEUR-WS, 2010.
- [8] C. Keßler, K. Janowicz, Linking Sensor Data - Why, to What, and How?, in: 3rd International Workshop on Semantic Sensor Networks, CEUR-WS, 2010.
- [9] W3C OWL Working Group, OWL 2 Web Ontology Language, W3C Recommendation, available at <http://www.w3.org/TR/owl-overview/> (last accessed 11th May 2012), 2009.
- [10] M. Compton, C. Henson, H. Neuhaus, L. Lefort, A. Sheth, A Survey of the Semantic Specification of Sensors, in: 2nd International Workshop on Semantic Sensor Networks, CEUR-WS, 17–32, 2009.
- [11] K. Janowicz, M. Compton, The Stimulus-Sensor-Observation Ontology Design Pattern and its Integration into the Semantic Sensor Network Ontology, in: 3rd International workshop on Semantic Sensor Networks, vol. 668, CEUR-WS, 2010.
- [12] N. Guarino, Formal Ontology and Information Systems, in: International Conference on Formal Ontology in Information Systems, IOS Press, 3–15, 1998.
- [13] A. Gangemi, Ontology Design Patterns for Semantic Web Content, in: Y. Gil, E. Motta, R. Benjamins, M. Musen (Eds.), 4th International Semantic Web Conference (ISWC 2005), vol. 3729 of *Lecture Notes in Computer Science*, Springer, 262–276, 2005.
- [14] RDF Vocabulary Description Language 1.0: RDF Schema, W3C Recommendation, available at <http://www.w3.org/TR/rdf-schema/> (last accessed 11th May 2012), 2010.
- [15] SKOS Simple Knowledge Organization System Reference, W3C Recommendation, available at <http://www.w3.org/TR/skos-reference> (last accessed 11th May 2012), 2009.
- [16] Joint Committee for Guides in Metrology (Working Group 2), International Vocabulary of Metrology — Basic and General Concepts and Associated Terms, JCGM 200:2008, 2008.

¹⁹<http://sweet.jpl.nasa.gov/> (last accessed 11th May 2012)

²⁰<http://www.sensorgrid4env.eu/> (last accessed 11th May 2012)

²¹<http://www.sensorgrid4env.eu/ontologies/> (last accessed 11th May 2012)

²²<http://www.ordnancesurvey.co.uk/oswebsite/ontology/> (last accessed 11th May 2012)

²³<http://aemet.linkeddata.es/> (last accessed 11th May 2012)

²⁴<http://www.ict-exalted.eu/> (last accessed 11th May 2012)

²⁵<http://www.csiro.au/science/Sensors-and-network-technologies.html> (last accessed 11th May 2012)

- [17] W. V. O. Quine, *From Stimulus to Science*, Harvard University Press, 1995.
- [18] W. Kuhn, A Functional Ontology of Observation and Measurement, in: *Proceedings of the 3rd International Conference on GeoSpatial Semantics*, Springer-Verlag, 26–43, 2009.
- [19] F. Probst, M. Lutz, Giving Meaning to GI Web Service Descriptions, in: *2nd International Workshop on Web Services: Modeling, Architecture and Infrastructure*, INSTICC Press, 23–35, 2004.
- [20] Mathematical Markup Language (MathML), W3C Recommendation, available at <http://www.w3.org/TR/MathML/> (last accessed 11th May 2012), 2010.
- [21] M. Compton, H. Neuhaus, K.-N. Tran, Reasoning about Sensors and Compositions, in: *2nd International Semantic Sensor Networks Workshop*, CEUR-WS, 33–48, 2009.
- [22] K. Taylor, P. Penkala, Using Explicit Semantic Representations for User Programming of Sensor Devices, in: *5th Australasian Ontology Workshop*, ACS, 47–55, 2009.
- [23] A. Moraru, D. Mladenic, M. Vucnik, M. Porcius, C. Fortuna, M. Mohoric, Exposing real world information for the web of things, in: *8th International Workshop on Information Integration on the Web*, ACM, 2011.
- [24] K. Wenzel, J. Riegel, A. Schlegel, M. Putz, Semantic Web Based Dynamic Energy Analysis and Forecasts in Manufacturing Engineering, in: *Proceedings of the 18th CIRP International Conference on Life Cycle Engineering*, Springer-Verlag, 507–512, 2011.
- [25] D. Corsar, P. Edwards, N. Velaga, J. Nelson, J. Pan, Short Paper: Addressing the Challenges of Semantic Citizen-Sensing, in: *4th International Workshop on Semantic Sensor Networks*, CEUR-WS, 90–95, 2011.
- [26] L. Yu, Y. Liu, Using the Linked Data Approach in a Heterogeneous Sensor Web: Challenges, Experiments and Lessons Learned, in: *Workshop on Sensor Web Enablement*, 2011.
- [27] M. Leggieri, A. Passant, M. Hauswirth, A Contextualised Cognitive Perspective for Linked Sensor Data - Short paper, in: *3rd International Workshop on Semantic Sensor Networks*, CEUR-WS, 2010.
- [28] D. Pfisterer, K. Römer, D. Bimschas, O. Kleine, R. Mietz, C. Truong, H. Hasemann, A. Kröller, M. Pagel, M. Hauswirth, M. Karnstedt, M. Leggieri, A. Passant, R. Richardson, SPITFIRE: toward a semantic web of things, *IEEE Communications Magazine* 49 (11) (2011) 40–48.
- [29] K. Janowicz, A. Bröring, C. Stasch, S. Schade, T. Everding, A. Llayes, A RESTful Proxy and Data Model for Linked Sensor Data, *International Journal of Digital Earth* .
- [30] A. Bröring, P. Maué, K. Janowicz, D. Nüst, C. Malewski, Semantically-enabled Sensor Plug & Play for the Sensor Web, *Sensors* 11 (2011) 7568–7605.
- [31] A. J. G. Gray, R. García-Castro, K. Kyzirakos, M. Karpathiotakis, J.-P. Calbimonte, K. R. Page, J. Sadler, A. Frazer, I. Galpin, A. A. A. Fernandes, N. W. Paton, Ó. Corcho, M. Koubarakis, D. D. Roure, K. Martínez, A. Gómez-Pérez, A Semantically Enabled Service Architecture for Mashups over Streaming and Stored Data, in: *8th Extended Semantic Web Conference*, Springer-Verlag, 300–314, 2011.
- [32] A. J. G. Gray, J. Sadler, O. Kit, K. Kyzirakos, M. Karpathiotakis, J.-P. Calbimonte, K. R. Page, R. García-Castro, A. Frazer, I. Galpin, A. A. A. Fernandes, N. W. Paton, Ó. Corcho, M. Koubarakis, D. D. Roure, K. Martínez, A. Gómez-Pérez, A Semantic Sensor Web for Environmental Decision Support Applications, *Sensors* 11 (9) (2011) 8855–8887.