

Sensor Data Provenance: SSNO and PROV-O Together at Last

Michael Compton¹, David Corsar², and Kerry Taylor^{1,3}

¹ CSIRO Digital Productivity and Services, Canberra
firstname.lastname@csiro.au

² University of Aberdeen, Aberdeen, UK dcorsar@abdn.ac.uk

³ Australian National University

Abstract. This paper presents an alignment between the W3C Provenance Working Group's recommended ontology (PROV-O) and the W3C Semantic Sensor Networks Incubator Group's ontology (SSNO). The alignment views PROV-O as an upper ontology which is extended with SSNO concepts and properties. This allows representation of observation details and sensor deployments that are not possible in the SSNO alone, and gives a basis for alignment with Open Geospatial Consortium Observations & Measurements aligned ontologies. Further to the alignment, rules are presented that further constrain the interpretation of the aligned ontologies and provide a mechanism by which provenance information can be generated from SSN data thereby allowing modellers to take advantage of the new features. The benefit of the aligned ontologies is illustrated with an example of cross-domain provenance querying enabled by the alignment.

1 Introduction

Sensor deployments for smart cities, the Internet of Things, crowd sensing, and environmental research produce large volumes of data, which, when analysed either in real-time or using archives, have the potential to impact on virtually all aspects of society [2]. The Semantic Sensor Web can play an important role in realising these impacts by supporting the identification, selection and use of sensor data through expressive representations of sensor resources. This support can be extended by providing additional details, such as how the data was produced, any processing to which it was subjected, and who was involved in those steps, i.e. the data's provenance [2, 21]. Provenance is valuable here as it can be used to assist users with understanding, verifying, and assessing the data's quality and trustworthiness before use [20, 28, 17].

Due to the work of the W3C Semantic Sensor Networks (SSN) Incubator Group and W3C Provenance Working Group, the semantic web community now has OWL2 [29] ontologies designed to support the reuse and interoperability of both sensor and provenance information. The SSN ontology (SSNO) [5, 16] has been adopted as a defacto standard for the semantic specification of sensors, sensor devices, systems, processes, and observations. The SSNO is designed to

capture both historical details (for example, details of deployments and how observation values were produced) along with current information (for example, sensor capabilities). The PROV-O ontology [14] is a W3C recommendation for the representation and interchange of provenance information on the Web, where provenance is defined as information about the entities, activities, and people involved in the production of things - data, physical objects, etc. [20]. PROV-O builds on significant research by the database, workflow, and e-science communities (see [3, 19, 25] for relevant reviews).

This paper contributes an alignment between the SSNO and PROV-O models, including rules that both constrain the interpretation of the alignment and can be used for inferring provenance from sensor data. The alignment extends the expressive capability of the SSNO for recording observations and system deployments, enabling more comprehensive historical information to be described than is possible using SSNO alone. As such, the alignment also serves as a best practice guide for using SSNO and PROVO to represent the provenance of observations, observation values, and sensor system deployments. The alignment also enables the integration of SSNO data with other data expressed using PROV-O, and so makes SSNO data available to tools and applications capable of consuming PROV-O (e.g. for visualisations).

Paper Outline: First, the SSN (§2.1) and PROV (§2.2) ontologies are discussed. Then related work (§3) is reviewed. The alignment (§4) is then presented in terms of a central pattern (§4.1), aligning the Stimulus-Sensor-Observation pattern from the SSNO to PROV-O along with an alignment (§4.2) for the SSNO’s platforms and deployments model. Next, rules are given (§4.3) that further inform the alignment and can be used to produce data in the aligned ontology from SSNO data. An example (§5) illustrates the use of the alignment. The paper concludes (§6) with a discussion of the alignment and its features.

2 Background

The PROV and SSN ontologies are accessible from their respective namespaces:

`http://www.w3.org/ns/prov`, and
`http://purl.oclc.org/NET/ssnx/ssn`

Throughout, the namespaces for the PROV and SSN ontologies are abbreviated as ‘prov’ and ‘ssn’ respectively. Hence `ssn:Sensor` means the concept `http://purl.oclc.org/NET/ssnx/ssn#Sensor`. The SSN ontology uses the DOLCE Ultralite ontology, called DUL, as an upper ontology and its namespace, `http://www.loa-cnr.it/ontologies/DUL.owl`, is abbreviated as ‘dul’. Concepts and properties given without a namespace are those of the alignment discussed here.

2.1 SSN

The SSN ontology was designed to describe sensors: what is observed, how observations are made, the observations, and the qualities of the sensors and ob-

servations. Full descriptions of the SSNO are given in Compton et al. [5] and the incubator group’s final report [16].

The SSNO is built around the Stimulus-Sensor-Observation pattern [11] that describes the relationship between an observing `ssn:Sensor`, the `ssn:Property` measured, the real-world `ssn:Stimulus`, the `ssn:Sensing` procedure followed and the resultant `ssn:Observation`.

The SSNO expands on the central pattern to describe the `ssn:Accuracy`, `ssn:Frequency`, `ssn:Drift`, etc. of a sensor as its `ssn:MeasurementCapability`, to describe the `ssn:OperatingRange` and `ssn:SurvivalRange` of sensors, and to provide a skeleton structure for describing how a sensor may be attached to an `ssn:Platform` and used in an `ssn:Deployment`.

However, it is the Stimulus-Sensor-Observation pattern that forms the key part for the alignment to PROV-O.

2.2 PROV

The PROV-O recommendation is an OWL2 encoding of the PROV Data Model [20], which describes provenance in terms of relationships between three main types of concepts: `prov:Entity`, which represents (physical, digital, or other types of) things; `prov:Activity`, which occur over time and can use and/or generate entities; and `prov:Agent`, which are responsible for activities occurring, entities existing, or another agent’s activity [14].

Relationships between these concepts describe the influence one has had on another. These include that an activity `prov:used` and `prov:generated` entities, ascribing an entity to an agent (`prov:wasAttributedTo`), and an agent to an activity (`prov:wasAssociatedWith`). The nature of the influence can be defined using qualified relations to describe the `prov:Role` of the entity, agent, or activity. Qualified relations include: `prov:Usage`, which defines the role of an entity used by an activity; and `prov:Association`, which defines the role of an agent in an activity, along with any `prov:Plan` the agent was following during the activity.

3 Related Work

The provenance of sensor data has many uses, including: supporting the understanding and reuse of sensor data, including data that has been aggregated or otherwise processed [21]; ensuring the correct attribution of publicly available sensor data [4]; supporting users trace the involvement of sensor data in experiments for reproducibility purposes [10]; searching for, and identifying sensor data within data stores [15]; verifying data transmitted through nodes in a sensor network [24]; and supporting quality [6] and trustworthiness assessments [28]. Despite this, there are few published alignments between sensor and provenance ontologies, which are discussed below.

The Open Provenance Model (OPM) [18] is used by Lie et. al. [17] to record the provenance of virtual sensors within the Tupelo semantic content management system. OPM integrates the sensor registration and selection events with

the retrieval of raw data and model-based transformations applied to derive new data. In this system, data is modelled as the OPM equivalent of `prov:Entity` and actions as `prov:Activity`; no further alignments are described.

Patni et. al. [22] use the Provenir ontology [23] to capture and store the provenance of sensor data in their sensor management system. Alignments between Provenir and their sensor ontology are defined through a series of `rdfs:subClassOf` and `rdfs:subPropertyOf` relationships modelling sensor-specific provenance information. This includes modelling observation values as the Provenir equivalent of `prov:Entity` and the sampling time property as the equivalent of `prov:atTime`.

Stasch et. al. [26] describe their extension of the SSNO to represent aggregations of observations, and detail the use of the Provenance Vocabulary⁴ and OPM to record each aggregation's provenance. Observations and aggregations are modelled using the Provenance Vocabulary and OPM equivalents of `prov:Entity`, with aggregations being created by an aggregation activity that used observations. However, it is unclear if the remaining SSNO concepts have been aligned to a provenance model, or if these alignments have been explicitly defined in an ontology to enable their reuse.

In the context of a citizen sensing application, Corsar et. al. [6] define a partial alignment between the SSNO and PROV-O. The alignment is restricted to defining subclass relationships between the SSNO observation, sensor, and sensing concepts and the PROV-O entity, agent, and activity concepts respectively. The alignment is used to integrate SSNO data with other data via the PROV-O model to support quality assessment of observations from citizen sensors.

These various works illustrate a requirement to combine sensor and provenance models, and potential uses of the resulting model. However, to the best of our knowledge, there does not currently exist a comprehensive alignment between any established sensor and provenance models. We have therefore chosen to develop such an alignment between the two ontologies (SSNO and PROV-O) that are now accepted within (and beyond) their respective communities as the main reference models.

4 SSNO-PROV-O Alignment

While the SSN and PROV ontologies are compatible in some areas, the two are modelled from different perspectives. Largely, the SSN ontology is about properties and potential: what sensors measure, how they measure, and the qualities of such measurements. While, on the other hand, the PROV ontology models what has occurred and how things were made: what the entities are, what produced them and how. The potential for overlap and alignment between the two ontologies is observations. Observations are the things that are produced by sensors and, thus from a provenance perspective, this production or generation

⁴ <http://trdf.sourceforge.net/provenance/ns.html>

is the key point in linking the two ontologies. Indeed, it is observations around which the following alignment is built.⁵

The PROV ontology is the more abstract of the two and thus the alignment places SSNO concepts and relations into the PROV-O hierarchy, making them subconcepts and subproperties of PROV-O concepts and properties. The PROV-O ontology is used like an upper ontology in this respect. This approach allows, for example, other provenance data and modelling to be used in conjunction with the sensor data in a provenance setting.

New concepts extending the PROV-O hierarchy are created to further glue the two ontologies together, as simply placing SSNO concepts into the PROV-O hierarchy does not complete the full richness of the alignment.

In making the alignment, modelling choices are present even at the initial stages and each choice has far reaching consequences for how other aspects are aligned. The placement of `ssn:Observation` and `ssn:Sensor` are the most central.

A choice was made in the SSN ontology to make an `ssn:Observation` a `dul:Situation` (i.e.: `ssn:Observation` \sqsubseteq `dul:Situation`). That is, an observation is an interpretation of real-world events and the results of those events: for example, a stimulus (wind) spins the cups on a wind sensor, generating a current, and through an equation modelling the relationship between this and the physical property of wind speed the sensor outputs a value that we can choose to interpret as an observation of the wind speed at that moment. The observation is the social construct of the interpretation, not the act of the observing itself.

On the other hand, the Open Geospatial Consortium (OGC) Observations and Measurements (O&M) (previously an OGC standard [7, 8], now an ISO standard [1]), sees an observation as an event: the event of sensing and producing the result. The SSNO and O&M attach essentially the same data to an observation — a value, a feature of interest, an observed property, etc. — but place the observation itself in a different context. The SSNO argues that O&M conflates two aspects of the observation: the act and the interpretation.

This dichotomy poses an immediate choice in the alignment. Following the SSNO model, `ssn:Observation` would be aligned with `prov:Entity` (`ssn:Observation` \sqsubseteq `prov:Entity`). But that choice reinforces the distinction between O&M and the SSNO. It may also make the alignment less useful as it would not fit with OGC models and would not, for example, be able to represent data from an O&M aligned ontology, such as that given by Cox [9]. Aligning to O&M would align `ssn:Observation` to `prov:Activity` (`ssn:Observation` \sqsubseteq `prov:Activity`). Such a choice might be passable in an SSN Ontology not aligned to DOLCE, but with the DOLCE alignment it would be problematic as an observation would be both a `dul:Object` and a `prov:Activity`. Since `dul:Object` is disjoint from `dul:Event` but PROV-O specifies no disjoints this alignment may not lead to inconsistency, but would be ontologically uncomfortable.

Instead of following either approach, the alignment reconciles these disparate viewpoints. It aligns the SSNO approach to PROV-O and describes new PROV-

⁵ The alignment ontology is available at <http://purl.oclc.org/NET/ssnprov/ssnprov>.

<i>ssn</i> : <i>Sensor</i>	\sqsubseteq <i>prov</i> : <i>Entity</i>	<i>ssn</i> : <i>Stimulus</i>	\sqsubseteq <i>prov</i> : <i>Activity</i>
	\sqsubseteq <i>prov</i> : <i>Agent</i>		
<i>ssn</i> : <i>ObservationValue</i>	\sqsubseteq <i>prov</i> : <i>Entity</i>	<i>ssn</i> : <i>SensorOutput</i>	\sqsubseteq <i>prov</i> : <i>Entity</i>
			\sqsubseteq <i>prov</i> : <i>alternateOf</i> some <i>ssn</i> : <i>ObservationValue</i>
<i>ssn</i> : <i>Sensing</i>	\sqsubseteq <i>prov</i> : <i>Plan</i>	<i>ssn</i> : <i>Observation</i>	\sqsubseteq <i>prov</i> : <i>Entity</i>
<i>ActivityOfSensing</i>	\sqsubseteq <i>prov</i> : <i>Activity</i>		
	\sqsubseteq <i>prov</i> : <i>generated</i> some <i>ssn</i> : <i>SensorOutput</i>		
	\sqsubseteq <i>prov</i> : <i>qualifiedAssociation</i> some <i>SensorPerformedSensing</i>		
	\sqsubseteq <i>prov</i> : <i>qualifiedUsage</i> some <i>SensingUsage</i>		
	\sqsubseteq <i>prov</i> : <i>wasInfluencedBy</i> some <i>ssn</i> : <i>Stimulus</i>		
<i>ObservationInterpretation</i>	\sqsubseteq <i>prov</i> : <i>Activity</i>		
	\sqsubseteq <i>prov</i> : <i>generated</i> some <i>ssn</i> : <i>Observation</i>		
	\sqsubseteq <i>prov</i> : <i>wasInformedBy</i> some <i>ActivityOfSensing</i>		
<i>SensorPerformedSensing</i>	\sqsubseteq <i>prov</i> : <i>Association</i>		
	\sqsubseteq <i>prov</i> : <i>hadPlan</i> some <i>ssn</i> : <i>Sensing</i>		
	\sqsubseteq <i>prov</i> : <i>hadRole</i> some <i>PerformedSensing</i>		
	\sqsubseteq <i>sensingAgent</i> some <i>ssn</i> : <i>Sensor</i>		
	$\text{sensingAgent} \circ \text{prov} : \text{hadPlan} \sqsubseteq \text{ssn} : \text{implements}$		
<i>SensingUsage</i>	\sqsubseteq <i>prov</i> : <i>Usage</i>	<i>PerformedSensing</i>	\sqsubseteq <i>prov</i> : <i>Role</i>
	\sqsubseteq <i>prov</i> : <i>entity</i> some <i>ssn</i> : <i>Sensing</i>	<i>SensingMethod</i>	\sqsubseteq <i>prov</i> : <i>Role</i>
	\sqsubseteq <i>prov</i> : <i>hadRole</i> some <i>SensingMethod</i>		
<i>sensingAgent</i>	\sqsubseteq <i>prov</i> : <i>agent</i>	<i>ssn</i> : <i>isProducedBy</i>	\sqsubseteq <i>prov</i> : <i>wasAttributedTo</i>
<i>ssn</i> : <i>observedBy</i>	\sqsubseteq <i>prov</i> : <i>wasAttributedTo</i>	<i>ssn</i> : <i>observationResult</i>	\sqsubseteq <i>prov</i> : <i>wasDerivedFrom</i>
<i>ssn</i> : <i>System</i>	\sqsubseteq <i>prov</i> : <i>Entity</i>	<i>ssn</i> : <i>Platform</i>	\sqsubseteq <i>prov</i> : <i>Entity</i>
	\sqsubseteq <i>prov</i> : <i>Collection</i>		\sqsubseteq <i>ssn</i> : <i>attachedSystem</i> only <i>ssn</i> : <i>System</i>
	\sqsubseteq <i>ssn</i> : <i>hasSubSystem</i> only <i>ssn</i> : <i>System</i>		\sqsubseteq <i>ssn</i> : <i>inDeployment</i> some <i>ssn</i> : <i>Deployment</i>
<i>ssn</i> : <i>DeploymentRelatedProcess</i>	\sqsubseteq <i>prov</i> : <i>Activity</i>		
	\sqsubseteq \forall <i>ssn</i> : <i>deploymentProcessPart</i> . <i>ssn</i> : <i>DeploymentRelatedProcess</i>		
<i>ssn</i> : <i>Deployment</i>	\sqsubseteq <i>ssn</i> : <i>DeploymentRelatedProcess</i>		
	\sqsubseteq <i>ssn</i> : <i>deployedOnPlatform</i> only <i>ssn</i> : <i>Platform</i>		
	\sqsubseteq \forall <i>ssn</i> : <i>deployedSystem</i> . <i>ssn</i> : <i>System</i>		
<i>ssn</i> : <i>deployedOnPlatform</i>	\sqsubseteq <i>prov</i> : <i>used</i>	<i>ssn</i> : <i>deployedSystem</i>	\sqsubseteq <i>prov</i> : <i>used</i>
<i>ssn</i> : <i>hasSubSystem</i>	\sqsubseteq <i>prov</i> : <i>hadMember</i>		

Fig. 2. The alignment of the SSNO to PROV-O. Grey indicates assertions from the SSNO provided here for reference.

Sensors and observations are linked in the SSNO by *ssn*:*madeObservation* and *ssn*:*observedBy*. The alignment adds *prov*:*wasAttributedTo*, i.e. that the observation entity was attributed to the sensor agent. Further linking the two are new concepts *SensorPerformedSensing*, *ActivityOfSensing* and *ObservationInterpretation*, which add the detail that describe the observation activity and fill in the O&M perspective as discussed in above.

The new concept *ActivityOfSensing* is the *prov*:*Activity* of performing the sensing. An *ActivityOfSensing* *prov*:*generated* the *ssn*:*SensorOutput*, was influenced by (*prov*:*wasInfluencedBy*) the *ssn*:*Stimulus* and, through *SensorPerformedSensing* (a *prov*:*Association*), *prov*:*wasAssociatedWith* the sensor. Further, an *ActivityOfSensing* may be specified as *prov*:*wasInformedBy* things such as the feature of interest or the observed property, though this isn't required in the alignment. It is the *ActivityOfSensing* that fills in the O&M perspective of observation.

The *prov*:*Activity* of *ObservationInterpretation* records the activity that interpreted the results of an *ActivityOfSensing* and resulted in (*prov*:*generated*) a *dul*:*Situation* that is the *ssn*:*Observation*. The activity of *ObservationInterpretation* *prov*:*wasInformedBy* the *ActivityOfSensing*. As with the *ActivityOfSensing*, the *ObservationInterpretation* may be *prov*:*wasInformedBy* some of the aspects recorded by the *ssn*:*Observation*.

Observations themselves can record in SSNO the `ssn:Stimulus` as a related event (`dul:includesEvent`) as well as a feature of interest and observed property. Here, as with `ActivityOfSensing` and `ObservationInterpretation`, the alignment is left open. These links could be added by specifying them as `prov:wasInfluencedBy` or `prov:wasInformedBy`, but we chose to leave the alignment flexible here and allow such links to be included, but not mandate them.

Specifying `ssn:observationResult` \sqsubseteq `prov:wasDerivedFrom` shows the provenance attribution of an observation as a situation being partly derived from the observation result.

The `ssn:SensorOutput` is `prov:wasAttributedTo` the sensor agent by virtue of specifying `ssn:isProducedBy` \sqsubseteq `prov:wasAttributedTo`. Some properties, of which `ssn:isProducedBy` is one, were not found to have alignments to DUL in the development of the SSNO. This is remedied in the alignment, which takes advantage of PROV-O as the upper ontology to further restrict the interpretation of `ssn:isProducedBy`.

The alignment adds further nuances to the description of an observation that SSN cannot do alone. `SensorPerformedSensing` can be enriched to show that the sensor had the role of performing the sensing in the `ActivityOfSensing` (`prov:hadRole` and `PerformedSensing`). More importantly, it can show the plan (`prov:hadPlan`) that was used to perform the sensing — linking the activity of sensing with the `ssn:Sensing` plan that was used. The SSNO alone can show what sensing plans a sensor is capable of performing, but, for an individual observation, the SSNO cannot show which plan was enacted. This advantage in the SSN-PROV-O alignment is helpful in specifications of the observations of multi-instruments and systems with complex sensing options, where the provenance can now record more accurately what was done.

The `dul:Plan` of `ssn:Sensing` in the SSN ontology can also be used to express if the observation was made in some particular way: i.e. specifying that a sensor `ssn:implements` some `ssn:Sensing` describes how the sensor works, while specifying that an observation had a `ssn:sensingMethodUsed` of some `ssn:Sensing` describes how the sensor was used in making the observation — a particular configuration or physical setup for example. The alignment enriches the picture by showing that it is the activity of sensing that used the plan. For this the alignment shows that an `ActivityOfSensing` can have a `prov:qualifiedUsage` of some `SensingUsage`.

The alignment further ensures that the `dul:Plan` a `ssn:Sensor` enacts in an `ActivityOfSensing` must be a plan that it `ssn:implements` by stating the role chain:

$$\text{sensingAgent}^- \circ \text{prov:hadPlan} \sqsubseteq \text{ssn:implements}.$$

There is no alignment to PROV-O for aspects of SSNO such as the measuring properties (`ssn:MeasurementCapability`) or `ssn:OperatingRange` as these are the static aspects of sensors that are covered by the SSNO and not PROV-O. That is, because aspects such as accuracy and drift, or specifications of operating and survival ranges are inherent properties of sensors, they have no natural alignment into PROV-O.

In fact, there is no alignment to PROV-O for any `ssn:Property` as properties such as accuracy (which is `ssn:MeasurementProperty`) or survival temperature (which would be `ssn:SurvivalProperty`) being measurable properties of sensors, just as temperature is a measurable property of a location. Properties could be aligned as `prov:Entity` (a conceptual entity in PROV-O), but since properties are inherent for an object and PROV-O doesn't provide mechanisms for talking about and linking entities, except through creation or participation in activities or membership of collections, there is no useful way to align properties to PROV-O. For example, one might want to describe the provenance of a sensor, including its creation and creator, but such a description is unlikely to involve specifying that the creator made or generated the accuracy of the sensor.

That there are aspects of the SSNO not covered by PROV-O and parts of the alignment that provide extra capability to the SSNO shows that the alignment extends each ontology.

A further alignment is possible for the deployments and platforms aspects of the SSNO as this describes time-varying aspects of sensors. Again the alignment adds expressive power not available in SSN alone.

4.2 Deployments Alignment

In the SSNO deployments are DUL processes (`ssn:Deployment` \sqsubseteq `dul:Process`): that is, events about which the evolving nature is important. Such a conception of a deployment as representing the ongoing process of initial deployment, maintenance, addition and removal, recalibration, etc. fits naturally with PROV-O. The alignment, however, adds to the expressive capability of the SSNO, and more nuanced and clear specifications can be made in the alignment, giving the expressive capability to state subtle properties of the evolution and nature of deployments and thus may give further reasoning power for sensor search and selection.

The alignment for deployments (bottom of Figure 2) makes the assertions that the processes for deployment are `prov:Activity`, that both `ssn:Platform` and `ssn:System` are `prov:Entity`, and that an `ssn:System` is a `prov:Collection`. Together with the properties assertions, the whole of the SSNO deployments skeleton is aligned to PROV-O.

Consider, for example, Figure 3 which shows an SSNO description of a deployment (black) augmented with a description that can be achieved in the aligned ontologies (grey). The SSNO can describe the deployment and the processes, systems and platforms involved, but can't show the relationships and derivations between the parts. Once the PROV-O parts are added, the usage, generation and temporal dependencies are clear, though one may wish to further specify these with timing assertions on the activities, such as with `prov:startedAtTime` and `prov:endedAtTime`, or by using `dul:precedes` and `dul:follows`.

If desired, information about the agents associated with the deployment-related processes can also be defined (not expressible in the SSNO). This, in turn, can include details explaining why the action took place: for example, as part of the system's maintenance plan or in response to the system malfunctioning.

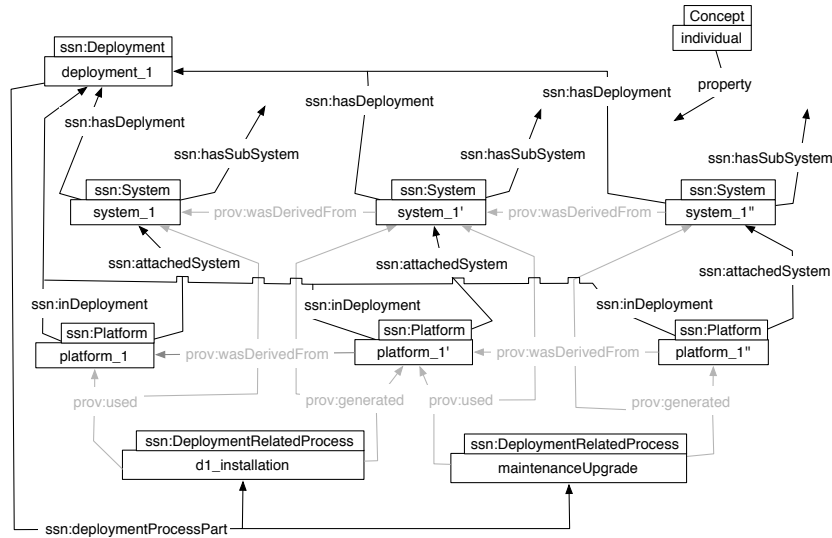


Fig. 3. An example of deployments. Black indicates SSNO only modelling; grey the extra modelling in the aligned ontology.

Further subtle nuances can be gained, for example, by adding that `system_1'`, `system_1''`, `platform_1'` and `platform_1''` from Figure 3 are revisions of `system_1` and `system_1`. This could be done in the aligned ontology by specifying `prov:wasRevisionOf`. A more complete specification could be gained by specifying temporal parts using a 4D fluents [30] approach.

Such detailed modelling could be used for example when sensor selection is dependent on dynamic properties of sensors, such as consistency of maintenance or time since last calibration. This information can also be used for identifying and selecting sensor data or to make judgements about the data from providers or networks based on the maintenance history of other deployments made by the provider. Further, calibration and maintenance events could be used, for example, in conjunction with a sensor's static specification of accuracy and drift to determine the likely current performance of the sensor based on its history. This, in turn, can be used during, for example, quality assessments of observations. Information about the agents involved in the deployment can be beneficial when making trust assessments [28], or for data access control (for example, restricting access to only members of the organisation responsible for the sensor).

4.3 Rules

As is often the case, the OWL specification restricts our interpretation to a set of potential satisfying models, but one may wish to add further detail. For this we

have defined a set of rules⁶ to further constrain the interpretation of the model and which can be used to generate aligned provenance data from SSN data.

The rules define how the concepts introduced by the alignment should be interpreted in the context of SSN data. Two types of rules could be considered: the first define ObservationInterpretation and ActivityOfSensing; the second define the relationships, particularly from ActivityOfSensing to the remaining concepts.

Implementations of these rules have been produced using the Stardog 2.1.3⁷ and SPARQL Inferencing Notation⁸ (SPIN) formats. The following shows fragments of our Stardog rules. Stardog rules⁹ are written in a SPARQL-like syntax with an ‘if ... then ...’ structure. The Stardog rules are activated on query and, although these infer the existence of new individuals and can use the inferred individuals in answering queries, the inferred individuals are not persistent. The SPIN rules, however, can persist the inferred individuals, but the SPIN reasoner is limited to using Jena¹⁰ compatible models.

The following fragment of a Stardog rule infers an ObservationInterpretation based on the existence of an ssn:Observation.¹¹

```
IF {
  ?obs rdf:type ssn:Observation .
  FILTER NOT EXISTS (?obsI prov:generated ?obs .
    ?obsI rdf:type prov:ssn:ObservationInterpretation) .
  BIND (UUID() AS ?obsInterp) .
} THEN { ?obsInterp prov:generated ?obs .
  ?obsInterp rdf:type prov:ssn:ObservationInterpretation .
}
```

The following example rule shows how the SensingUsage qualified association can be derived for an ActivityOfSensing.

```
IF {
  ?obs rdf:type ssn:Observation .
  ?obs ssn:sensingMethodUsed ?sensing .
  ?obsI prov:generated ?obs .
  ?obsI rdf:type prov:ssn:ObservationInterpretation .
  ?obsI prov:wasInformedBy ?actOfSen .
  ?actOfSen rdf:type prov:ssn:ActivityOfSensing .
  FILTER NOT EXISTS (?actOfSen prov:qualifiedUsage ?senUse .
    ?senUse rdf:type prov:ssn:SensingUsage) .
  BIND (UUID() AS ?sensingUsage) .
} THEN { ?actOfSen prov:qualifiedUsage ?sensingUsage .
  ?sensingUsage rdf:type prov:ssn:SensingUsage .
}
```

⁶ The rules are available at <http://purl.oclc.org/NET/ssnprov/rules>.

⁷ <http://stardog.com/>

⁸ <http://spinrdf.org/>

⁹ <http://docs.stardog.com/owl2/>

¹⁰ <https://jena.apache.org/>

¹¹ The UUID() function creates a unique identifier for the new individual.

```

    ?sensingUsage prov:entity ?sensing .
    ?sensingUsage prov:hadRole prov:ssn:SensingMethod .
}

```

5 Example

One of our motivations for this paper is to support the enrichment of PROV-O with sensor-specific concepts of relevance to provenance. Simultaneously, we enable PROV-O to act as a common language for modelling provenance-like interactions between the data produced by sensors and non-sensor data, such as that sourced from social networks or simulation systems. In this example we validate those aims by demonstrating the interaction of our alignment with the conceptual model *CERIF 1.3*. We show queries over provenance independently recorded as SSNO or CERIF but jointly queryable (a) through PROV-O by virtue of the alignment mappings and basic OWL inference and (b) through the extended terminology of both SSNO and CERIF for domain-specific precision.

CERIF 1.3 is natively described as a large relational data model that supports the management of research information, associating comprehensive information about European research projects and infrastructure with their resources and products [12]. The model includes sophisticated support for encoding ontological relationships between the concepts generally represented as tables in the model. A comprehensive re-interpretation of the model as an OWL ontology could both unify the “semantics” represented in the relationships with the semantics represented as relational attributes, and could also improve the interoperability of data using linked data approaches. An initial OWL re-interpretation described in [13] models a few of the relationships but does not cover any attributes.

CERIF describes concepts that could be mapped to PROV-O and thereby enrich the provenance of research results, being patents, publications or products such as datasets and software. CERIF links these results to equipment, funding, people, and impact indicators. We have developed a partial mapping to PROV-O based on our own partial encoding of CERIF as an OWL 2 ontology, which enables us to demonstrate cross-cutting SPARQL queries that relate information arising from the domain of SSNO to information in the domain of CERIF, using PROV-O as a lingua-franca while using each of those ontologies for greater specificity when required.

For the queries, let us assume a scenario where we have many sensors installed on an experimental farm that is used as a research facility by many independent research projects. The sensor network is described by SSNO, so the provenance of research results can be tracked back to data sources. In fact, this scenario has been realized [27]. Let us further assume that CERIF (with namespace prefix ‘cf’) is used for those research projects, and that the mapping described above is deployed. Now, a paddock inspection reveals that a post, kirbypost23 on which several sensing devices were mounted, has fallen down. At least some of the sensors appear to be continuing to operate, but they would be unreliable, especially the automatic weather station wxt520 that needs to be

both vertical, unobstructed and at a known elevation above the ground for wind and precipitation measurements. We want to contact the principal investigators of research projects that are using those sensors to let them know that the data might be of poor quality.

In the first query we use PROV-O terms exclusively to retrieve individuals that are encoded as instances of either SSNO or CERIF classes, but which are related through PROV-O by virtue of the respective alignments. By this we demonstrate the lingua-franca value of the alignment. The query retrieves all projects (activities) and their associated people (agents) that either used the weather station or used the sensor network that included the weather station.

```
SELECT ?proj ?sys ?pers
WHERE {
  { ?proj prov:used ?sys . ?sys prov:hadMember :wxt520 }
  UNION
  { ?proj prov:used :wxt520 }
}.
?proj prov:wasAssociatedWith ?pers .
?pers a prov:Person }
```

In this query, constrained to PROV-O terms, we were unable to look for *all* the sensors installed on kirbypost23, *only* the principal investigators of the projects, or *only* those projects where the sensor was used for observing (as opposed to being a photographic subject, for example). Our second query is better targeted, handling these issues by employing the more specific modelling available in both of SSNO and CERIF. It demonstrates enrichment of PROV-O with both sensor-specific and research management-specific concepts of relevance to provenance, through independent alignments of each to PROV-O.

```
SELECT ?proj ?sys ?device ?pers
WHERE {
  ?proj cf:ObservedWith ?sys .
  ?proj cf:PrincipalInvestigator ?pers .
  {
    { ?sys dul:hasPart ?device .
      ?device a ssn:SensingDevice .
      ?device ssn:onPlatform :kirbypost23
    }
    UNION
    { ?sys a ssn:SensingDevice .
      ?sys ssn:onPlatform :kirbypost23
    }
  }
}}
```

These are rather simple queries aimed to compactly demonstrate the value of the mapping. Clearly much more could be retrieved by queries utilising more of the terminology available in PROV-O, SSNO, and CERIF, such as the time period during which the post would have fallen.

6 Conclusion

This paper presented an alignment between the W3C recommendation PROV-O and the current defacto standard for the semantic description of sensors, the SSNO. As well as aligning SSNO concepts and properties to PROV-O, further detail was gained by creating new concepts in the PROV-O hierarchy and linking them to both the SSNO and PROV-O. The alignment links SSNO-based ontologies and observational data to provenance, and is capable of more detailed modelling for sensors than the SSNO alone. In particular, the alignment extends the modelling of how the sensing took place and provides capabilities for detailed modelling of the passage of time in relation to deployments and the changes that take place in installation, maintenance, and upgrade. The extra detail for observation descriptions aligns the SSNO to the O&M view, drawing the two into the same framework and allowing extra interoperability.

Interestingly, not all of the SSNO is aligned to PROV-O. The alignment reflects that provenance describes what has happened, and, hence, PROV-O is not strong on entity to entity relations, meaning that the parts of the SSNO that describe properties of sensors are left unmapped. However, these unmapped properties can still be used to advantage in a provenance context as demonstrated by our second query example. The alignment doesn't align DUL and PROV-O. We felt that the useful alignment was SSNO to PROV-O and that a PROV-O to DUL alignment doesn't provide the same benefit and places restrictions on the meaning of PROV-O that may make the alignment less useful.

Rules provided as part of the alignment further constrain the interpretation of the relationship between the two ontologies. They guide users and implementing tools with using the alignment to define provenance for sensor data. Alternatively, the rules enable SSNO observational data to be automatically enriched with the extra concepts from the alignment and be used in a provenance context.

As sensors continue to become more ubiquitous and the availability of semantic sensor data increases, its use will become even more commonplace than it is today. By documenting data provenance, the alignment described in this paper can play an important role in supporting agents to understand and utilise such data. Given the status of the SSNO and PROV-O as reference models for sensor and provenance descriptions, we believe that this alignment will find service in both communities, and may also act as a guide to others who require to describe the provenance of sensed data.

Acknowledgements The authors would like to thank Keith Jeffery for assistance in interpreting the CERIF data model. David Corsar is supported by the award made by the RCUK Digital Economy programme to the dot.rural Digital Economy Hub (award reference EP/G066051/1).

References

1. Geographic Information – Observations and Measurements. International Standard ISO 19156, International Organization for Standardization (2011)

2. Bisdikian, C., Kaplan, L., Srivastava, M.: On the quality and value of information in sensor networks. *ACM Transactions on Sensor Networks* 9(4), 48:1–48:26 (Jul 2013)
3. Cheney, J., Chiticariu, L., Tan, W.C.: Provenance in databases: Why, how, and where. *Trends databases* 1(4), 379–474 (April 2009)
4. Chong, S., Skalka, C., Vaughan, J.: Self-identifying sensor data. In: *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*. pp. 82–93. *IPSN '10*, ACM, New York, NY, USA (2010)
5. Compton, M., Barnaghi, P., Bermudez, L., Garcia-Castro, R., Corcho, O., Cox, S., Graybeal, J., Hauswirth, M., Henson, C., Herzog, A., Huang, V., Janowicz, K., Kelsey, W., Phuoc, D.L., Lefort, L., Leggieri, M., Neuhaus, H., Nikolov, A., Page, K., Passant, A., Sheth, A., Taylor, K.: The SSN ontology of the W3C semantic sensor network incubator group. *Journal of Web Semantics: Science, Services and Agents on the World Wide Web* 17 (2012)
6. Corsar, D., Edwards, P., Baillie, C., Markovic, M., Papangelis, K., Nelson, J.: Short paper: Citizen sensing within a real time passenger information system. In: *Proceedings of the 6th International Workshop on Semantic Sensor Networks*. vol. 1063, pp. 77–82. *CEUR Workshop Proceedings* (2013)
7. Cox, S.: Observations and Measurements — Part 1 — Observation schema. *OpenGIS Implementation Standard OGC 07-022r1*, Open Geospatial Consortium Inc. (December 2007)
8. Cox, S.: Observations and Measurements — Part 2 — Sampling Features. *OpenGIS Implementation Standard OGC 07-002r3*, Open Geospatial Consortium Inc. (December 2007)
9. Cox, S.: An explicit OWL representation of ISO/OGC Observations and Measurements. In: *6th International workshop on Semantic Sensor Networks*. vol. 1063. *CEUR-WS* (2013)
10. Hensley, Z., Sanyal, J., New, J.: Provenance in sensor data management. *Queue* 11(12), 50:50–50:63 (December 2013)
11. Janowicz, K., Compton, M.: 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. Jorg, B., Jeffery, K., Dvorak, J., *et. al.*: CERIF 1.3 full data model (FDM) introduction and specification. Technical report, euroCRIS (January 2012)
13. Jorg, B., Lappalainen, J., Kastrantas, K.: Modeling the semantics of contextual and content-specific research metadata using ontology languages: issues on combining CERIF and OWL. In: *International Conference on Computational Science, ICCS 2012*. Elsevier (2012)
14. Lebo, T., Sahoo, S., McGuinness, D., Belhajjame, K., Cheney, J., Corsar, D., Garijo, D., Soiland-Reyes, S., Zednik, S., Zhao, J.: PROV-O: The PROV ontology. *W3C Recommendation* (2013), available at <http://www.w3.org/TR/prov-o/> (last accessed 23rd April 2014)
15. Ledlie, J., Ng, C., Holland, D., Muniswamy-Reddy, K., Braun, U., Seltze, M.: Provenance-aware sensor data storage. In: *Proceedings of the 1st IEEE International Workshop on Networking Meets Databases (NetDB)* (April 2005)
16. Lefort, L., Henson, C., Taylor, K., Barnaghi, P., Compton, M., Corcho, O., Castro, R.G., Graybeal, J., Herzog, A., Janowicz, K., Neuhaus, H., Nikolov, A., Page, K.: *Semantic Sensor Network XG Final Report*. *W3C Incubator Report* (2011), available at <http://www.w3.org/2005/Incubator/ssn/XGR-ssn-20110628/> (last accessed 23rd April 2014)

17. Liu, Y., Futrelle, J., Myers, J., Rodriguez, A., Kooper, R.: A provenance-aware virtual sensor system using the open provenance model. In: Collaborative Technologies and Systems (CTS), 2010 International Symposium on. pp. 330–339 (May 2010)
18. Moreau, L., Clifford, B., Freire, J., Futrelle, J., Gil, Y., Groth, P., Kwasnikowska, N., Miles, S., Missier, P., Myers, J., Plale, B., Simmhan, Y., Stephan, E., Van den Bussche, J.: The open provenance model core specification (v1.1). *Future Generation Computer Systems* 27(6), 743–756 (June 2011)
19. Moreau, L., Ludäscher, D.: Special issue: The first provenance challenge. *Concurrency and Computation: Practice and Experience* 20(5), 409–418 (April 2008)
20. Moreau, L., Missier, P., Belhajjame, K., B’Far, R., Cheney, J., Coppens, S., Cresswell, S., Gil, Y., Groth, P., Klyne, G., Lebo, T., McCusker, J., Miles, S., Myers, J., Sahoo, S., Tilmes, C.: PROV-DM: The PROV Data Model. W3C Recommendation (2013), available at <http://www.w3.org/TR/prov-dm/> (last accessed 23rd April 2014)
21. Park, U., Heidemann, J.: Provenance in sensor network republishing. In: Freire, J., Koop, D., Moreau, L. (eds.) *Provenance and Annotation of Data and Processes*, Lecture Notes in Computer Science, vol. 5272, pp. 280–292. Springer Berlin Heidelberg (2008)
22. Patni, H., Sahoo, S., Henson, C., Sheth, A.: Provenance aware linked sensor data. In: *Proceedings of the 2nd Workshop on Trust and Privacy on the Social and Semantic Web*. vol. 5 (May 2010)
23. Sahoo, S., Barga, R., Goldstein, J., Sheth, A., Thirunarayan, K.: Where did you come from...where did you go? An algebra and RDF query engine for provenance. Tech. rep., Kno.e.sis Center, Wright State University (2009)
24. Shebaro, B., Sultana, S., Gopavaram, S., Bertino, E.: Demonstrating a lightweight data provenance for sensor networks. In: Yu, T., Danezis, G., Gligor, V.D. (eds.) *Proceedings of the 2012 ACM Conference on Computer and Communications Security*. pp. 1022–1024. ACM (2012)
25. Simmhan, Y., Plale, B., Gannon, D.: A survey of data provenance in e-science. *SIGMOD Record* 34, 31–36 (2005)
26. Stasch, C., Sven Schad, S., Alejandro Llaves, L., Krzysztof Janowicz, K., Broring, A.: Aggregating linked sensor data. In: Taylor, K., Ayyagari, A., De Roure, D. (eds.) *Proceedings of the 4th International Workshop on Semantic Sensor Networks*. vol. 839, pp. 55–68 (October 2011)
27. Taylor, K., Lamb, D., Griffith, C., Falzon, G., Lefort, L., Gaire, R., Compton, M., Trotter, M., Wark, T.: Farming the web of things. *IEEE Intelligent Systems* 28(6), 12–19 (November-December 2013)
28. Umuhoza, D., Braun, R.: Trustworthiness assessment of knowledge on the semantic sensor web by provenance integration. In: *Proceedings of the 26th International Conference on Advanced Information Networking and Applications Workshops (WAINA)*. pp. 387–392 (March 2012)
29. W3C OWL Working Group: OWL 2 Web Ontology Language. W3C Recommendation (2009), available at <http://www.w3.org/TR/owl-overview/> (last accessed 23rd April 2014)
30. Welty, C., Fikes, R.: A reusable ontology for fluents in OWL. In: *Proceedings of the Fourth International Conference on Formal Ontology in Information Systems (FOIS)* (2006)