



SEVENTH FRAMEWORK PROGRAMME

Specific Targeted Research Project

Project Number:	FP7–SMARTCITIES–2013(ICT)
Project Acronym:	VITAL
Project Title:	Virtualized programmable InTerFACES for innovative cost-effective IoT depLoyments in smart cities

**D3.1.1 Virtual Models,
Data and Metadata for ICOs V1**

Document Id:	VITAL-D311-010914-draft
File Name:	VITAL-D311-010914-draft.docx
Document reference:	Deliverable 3.1.1
Version:	Draft
Editors:	Gregor Schiele, Thomas Geraghty
Organisation:	NUI Galway
Date:	01 / 09 / 2014
Document type:	Deliverable
Security:	PU (Public)

Copyright © 2013, 2014 VITAL Consortium

PROPRIETARY RIGHTS STATEMENT

This document contains information, which is proprietary to the VITAL Consortium. Neither this document nor the information contained herein shall be used, duplicated or communicated by any means to any third party, in whole or in parts, except with prior written consent of the consortium

DOCUMENT HISTORY

Rev.	Author(s)	Organisation(s)	Date	Comments
V001	Gregor Schiele	NUI Galway	23/07/2014	Initial version of the document
V002	Gregor Schiele	NUI Galway	28/07/2014	Added introduction
V003	Gregor Schiele	NUI Galway	07/08/2014	Additions on sensor metadata and measurements
V004	Gregor Schiele	NUI Galway	18/08/2014	Web and IoT services, CEP
V005	Anne Helmreich	NUI Galway	19/08/2014	Linked Data, events
V006	Miguel Mateo Montero	ATOS	19/08/2014	CEP
V007	Thomas Geraghty	NUI Galway	19/08/2014	Smart Cities, access control
V008	Gregor Schiele	NUI Galway	20/08/2014	VITAL IoT services, locations, fixes
V009	Thomas Geraghty	NUI Galway	21/08/2014	Updated Access Control
V010	Gregor Schiele, Anne Helmreich	NUI Galway	21/08/2014	Many updates in all parts, e.g. VITAL system, trust, JSON-LD, trust, etc.
V011	Gregor Schiele	NUI Galway	24/08/2014	Lots of polishing for existing parts, restructuring, IoT systems, linked data rewrite, data access service
V012	Thomas Geraghty	NUI Galway	25/08/2014	Added Unit of Measurement
V013	Gregor Schiele	NUI Galway	25/08/2014	Added discovery and filtering, rewrite sensor part slightly to include DCO, trust fixed, CEP fixed, monitoring added
V014	Thomas Geraghty	NUI Galway	26/08/2014	Added User and Groups model and updated integrated authorization and access control strategy
V015	Thomas Geraghty	NUI Galway	27/08/2014	Added first part of smart transport
V016	Thomas Geraghty	NUI Galway	28/08/2014	Added smart traffic services and conclusion to Smart Transport
V017	Gregor Schiele	NUI Galway	28/08/2014	Polishing and minor fixes, version ready for final reviews
V018	Angelos Lenios	Singularlogic	29/08/2014	Technical & Quality Review
V019	John Soldatos	AIT	30/08/2014	Technical & Quality Review
V020	Thomas Geraghty	NUI Galway	30/08/2014	Changes based on reviews, candidate for draft
Draft	Thomas Geraghty	NUI Galway	30/08/2014	Ready for EC submission
Draft	Gregor Schiele	NUI Galway	01/09/2014	Minor fixes, ready for EC submission

TABLE OF CONTENTS

1	INTRODUCTION	6
2	LINKED DATA AND SEMANTIC DESCRIPTIONS	7
2.1	RDF.....	7
2.2	JSON-LD.....	8
2.3	SPARQL	9
2.4	ONTOLOGIES	10
2.5	CONCLUSION	10
3	SENSORS AND MEASUREMENTS	12
3.1	THE W3C SEMANTIC SENSOR NETWORK ONTOLOGY.....	12
3.2	BASIC CONCEPTS AND ONTOLOGIES	14
3.3	SENSORS	17
3.4	SENSOR MEASUREMENTS.....	20
3.5	CONCLUSION	21
4	SMART CITIES.....	22
4.1	CITIES	22
4.2	SMART TRANSPORT	23
4.3	CONCLUSION	24
5	IOT SYSTEMS AND SERVICES.....	25
5.1	IOT SYSTEMS.....	25
5.2	IOT SERVICES.....	29
5.3	CONCLUSION	34
6	VITAL SYSTEM AND SERVICES.....	36
6.1	VITAL SYSTEM.....	36
6.2	DATA ACCESS.....	38
6.3	SECURITY AND ACCESS CONTROL.....	38
6.4	TRUST	42
6.5	DISCOVERY AND FILTERING	43
6.6	COMPLEX EVENT PROCESSING	48
6.7	MONITORING.....	49
6.8	ORCHESTRATION AND DEVELOPMENT.....	51
6.9	CONCLUSION	51
7	CONCLUSION AND NEXT STEPS.....	52
8	REFERENCES	53

LIST OF FIGURES

FIGURE 1: SPARQL EXAMPLE [LWS13]	10
FIGURE 2: EXAMPLE TIME INSTANT SPECIFICATION WITH OWL TIME	15
FIGURE 3: EXAMPLE WGS84 LOCATION WITH THE BASIC GEO VOCABULARY	15
FIGURE 4: JSON-LD CONTEXT FOR SENSOR.....	18
FIGURE 5: EXAMPLE SENSOR DESCRIPTION	19
FIGURE 6: EXAMPLE MEASUREMENT.....	20
FIGURE 7: JSON-LD CONTEXT FOR MEASUREMENT	21
FIGURE 8: EXAMPLE IOT SYSTEM DESCRIPTION	26
FIGURE 9: JSON-LD CONTEXT SPECIFICATION FOR SYSTEMS.....	27
FIGURE 10: IOT SYSTEM DESCRIPTION WITH PROVIDED SERVICES.....	29
FIGURE 11: EXAMPLE IOT SERVICE FOR PARKING SPACE MANAGEMENT (N3)	32
FIGURE 12: EXAMPLE IOT SERVICE FOR PARKING SPACE MANAGEMENT (JSON-LD)	33
FIGURE 13: JSON-LD CONTEXT SPECIFICATION FOR SERVICES.....	34
FIGURE 14: EXAMPLE VITAL DEPLOYMENT DESCRIPTION	37
FIGURE 15: EXAMPLES OF USERS AND GROUPS IN VITAL	40
FIGURE 16: ACCESS CONTROL POLICY FOR SINGLE SENSOR'S OBSERVATIONS	41
FIGURE 17: EXAMPLE SENSOR DESCRIPTION WITH PREDICTED MOVEMENT PATTERN	45
FIGURE 18: EXAMPLE SENSOR DESCRIPTION WITH NETWORK CONNECTION DATA	46
FIGURE 19: EXAMPLE SENSOR DESCRIPTION WITH LOCALIZER.....	48
FIGURE 20: EXAMPLE VIRTUAL MEASUREMENT FOR SYSTEM UPTIME.....	50
FIGURE 21: EXAMPLE SENSOR DESCRIPTION WITH HARDWARE COMPONENTS.....	51

LIST OF TABLES

TABLE 1: ONTOLOGY/ LANGUAGE PREFIXES.....	11
TABLE 2: FIELDS FOR MODELLING USERS AND GROUPS	39

TERMS AND ACRONYMS

ACL	Access Control List
API	Application programming interface
CEP	Complex Event Processing
COMANTO	Context Management Ontology
COMPOSE	Collaborative Open Market to Place Objects at your Service
CQELS	Continuous Query Evaluation over Linked Stream
CRUD	Create Read Update Delete
DCMI	Dublin Core Metadata Initiative
DCO	Delivery Context Ontology
DOLCE-Lite	Descriptive Ontology for Linguistic and Cognitive Engineering Lite
GEO	Basic Geo Ontology
GPS	Global Positioning System
GSN	Global Sensor Network
hRest	HTML for REST
HTTP	Hypertext Transfer Protocol
ICO	Internet Connected Object
ICT	Information and communication technology
IoT	Internet of Things
IRI	Internationalised Resource Identifier
JSON	JavaScript Object Notation
JSON-LD	JSON for Linked Data
LODE	Ontology for Linking Open Descriptions of Events
MSM	Minimal Service Model
MUO	Measurements Unit Ontology
N3	Notion 3
OAE	Ontology of Adverse Events
OM	Ontology of Units of Measure
OnTraJaCS	Ontology Based Traffic Control Systems
OWL	Web Ontology Language
OWL-S	Web Ontology Language for Web Services
PPI	Platform Provider Interface
PPO	Privacy Preference Ontology
QUDT	Quantities, Units, Dimensions and Data Types Ontologies
RDF	Resource Description Framework
REST	Representational State Transfer

S4AC	Social Semantic SPARQL Security for Access Control
SA-REST	Semantic Annotations for REST
SAWSDL	Semantically Annotated WSDL
SOAP	Simple Object Access Protocol
SOUPA	Standard Ontology for Ubiquitous and Pervasive Applications
SPARQL	SPARQL Protocol and RDF Query Language
SSN	Semantic Sensor Network
SSO	Stimulus-Sensor-Observation pattern
TAC	Triple Access Control
TAO	Trust Assertion Ontology
UO	Units of Measurement Ontology
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
UTO	Urban Traffic Ontology
W3C	World Wide Web Consortium
WAC	Web Access Control
WADL	Web API Definition Language
WSDL	Web Service Definition Language
WSMO	Web Service Modelling Ontology

1 INTRODUCTION

This deliverable specifies the data models and ontologies used in the VITAL project. These models and ontologies are used in three different ways. First, they are used by applications to communicate data between them and the VITAL platform, e.g. querying sensor data. Second, they are used by the so-called platform provider interface (PPI), which specifies an interface between VITAL and external IoT systems that should be integrated into VITAL, e.g. a description of the services offered by the IoT system. Third, they are used inside the VITAL platform to communicate data between VITAL services, e.g. security level information, etc.

Due to this diverse set of users (and usage cases), VITAL has to cover a wide area of data models. It has to specify how basic IoT sensors and sensor measurements are modelled, e.g. using the Semantic Sensor Network (SSN) ontology. It has to define how to describe IoT systems and IoT services, e.g. using the Minimal Service Model (MSM). It has to identify data models for Smart City applications, e.g. for smart transport systems. And it has to specify metadata for the VITAL system and its components, e.g. to model security and monitoring information.

In this document we analyse existing ontologies and data models that can be used as the basis of the VITAL data models and extend them as needed by VITAL. We subdivide the work into different areas – sensors and sensor measurements, Smart Cities, IoT systems and services, VITAL systems and services. For each of these areas we discuss which ontologies are used for modelling the required data (and why), present necessary additional data items and how to model them, and show examples of the resulting descriptions. The presented data models are not final and/or static. Instead, in the time until the second version of this deliverable (D3.1.2 in M24) we will fine-tune the data models, incorporate practical experiences and new requirements, and integrate new data items that we identified during the project lifetime.

The document is intended for: (a) users of the VITAL platform that want to learn about the used data models to develop applications, (b) IoT platform providers that want to learn how to integrate their platform with VITAL, (c) the partners of the VITAL consortium, allowing them to develop their components in a way that ensures easy integration and interoperability, (d) external researchers and developers that want to design and/or use data models in Smart City systems beyond VITAL, as well as (e) the project reviewers to better understand the work done in the project.

The document is structured as follows. First we give a brief overview on semantic and linked data technologies. Then we present the different ontologies – as well as necessary extensions – for the VITAL data model, e.g. for modelling sensor measurements and sensor descriptions, for Smart Cities, for IoT systems and services (including the data required for the PPI), as well as for the VITAL system itself and its services. We finish the deliverable with a short conclusion and outlook on further work during the VITAL project and beyond.

2 LINKED DATA AND SEMANTIC DESCRIPTIONS

This chapter provides background information required to understand the remainder of this deliverable. We give a short overview on linked data and its technologies. First, we explain the basic concepts of Linked Data. After that we briefly present the most commonly used technologies of Linked Data, namely RDF, SPARQL and Ontologies. Finally, we discuss how Linked Data will be used in VITAL.

The term ‘Linked Data’ is usually applied to a set of techniques for publishing and interlinking structured data on the Web. With Linked Data different kinds of data sources can be integrated (see [HeBi11]). Linked Data is based on four principles, explained in [HeBi11] and [HKHD09]:

1. The first principle is to **use URIs as names for things**. After the identification of items in a domain of interest, they are described in the data set with their properties and relationships. Each thing must be assigned a globally unique name, usually an HTTP URI (since this makes it easy to enforce global uniqueness). The fact that each thing has a globally unique name makes it possible to integrate local data sets with remote sets that have been developed independently.
2. The second principle is to **use HTTP URIs**, so clients can retrieve a description of the names thing or resource using the HTTP protocol. For humans the description can be provided as HTML, for machines it can be provided as RDF triples.
3. The third principle is to **use standards to provide information**. This usually means that standards like RDF or SPARQL are used to model and access data. We explain the most important standards later in this chapter.
4. The last principle is to **link to other URIs** and to enable the possibility to discover more things. That means that there should be external links pointing to other data sources on the Web. By following these links, a larger (distributed) data space can be explored automatically.

In a nutshell, Linked Data enables the implementation of generic applications operating over a huge, interconnected (distributed) data space by using Web standards and a common data model.

2.1 RDF

The most common data model used in the context of Linked Data is the Resource Description Framework (RDF) [GaSc14], which we introduce in this subsection. We first describe the basics of RDF, before discussing some of its advantages. RDF is a popular standard for describing things (known as resources or entities). By itself it is a graph-based data model that represents information as labelled directed graphs. This graph is built of triples that describe the data. Each triple (s, p, o) consists of a subject s , a predicate p and an object o . Take for example the information “The sky has the colour blue”. This would be modelled as a triple with “Sky” as the subject, “has colour” as the predicate and “blue” as the object. Note that the predicate in the middle always denotes the relationship between subject and object. Both the subject and the predicate are identified by URIs (assigning them globally unique identifiers)¹

¹ A subject can also be identified by a so-called blank node. A blank node is a local identifier.

while the object can be a URI or a literal value (i.e. a string or a number). The triple in our example could e.g. look like this (given in N-Triple notation):

```
<http://example.com/Sky> <http://example.com/hasColour> "blue"
```

Using RDF in a Linked Data context has some advantages. In the following we briefly outline the most important ones of these advantages. The reader is referred to [HeBi11] for more detailed information. Possible advantages are:

- If the identifiers of data items (both used as subjects and objects) and vocabulary terms (used as predicates) are HTTP URIs, the RDF data model can be used at global scale and anybody is able to refer to anything.
- Each RDF triple is included in the Web of Data and can be the starting point for explorations in the data space, because any URI can be looked up in an RDF graph over the Web.
- It is possible to set RDF links between data from different sources.
- Sets of triples can be merged in a single graph to combine information.
- Terms taken from different vocabularies can be mixed in a RDF graph.

As RDF is just a data model and not a data format, there are a number of data formats that can be used to write RDF data, either directly as triples or as nodes that can be mapped to RDF triples, e.g. RDF/XML [GaSc14], RDFa [HASB13], Turtle, N-Triples and JSON-LD. In the VITAL Project JSON-LD is used.

2.2 JSON-LD

Many developers have little or no experience with Linked Data, RDF or common RDF serialization formats such as N-Triples and Turtle. This produces extra overhead in the form of a steeper learning curve when integrating new systems to consume linked data. To counter this the project consortium decided to use a format based on a common serialization format such as XML or JSON. Thus, the two remaining options are RDF/XML and JSON-LD [SLK+14] [jsld]. JSON-LD was chosen over RDF/XML as the data format for all Linked Data items in VITAL. JSON-LD is a JSON-based serialisation for Linked Data with the following design goals:

- **Simplicity:** There is no need for extra processors or software libraries, just the knowledge of some basic keywords.
- **Compatibility:** JSON-LD documents are always valid JSON documents, so the standard libraries from JSON can be used.
- **Expressiveness:** Real-world data models can be expressed because the syntax serialises a directed graph.
- **Terseness:** The syntax is readable for humans and developers need little effort to use it.
- **Zero Edits:** Most of the time JSON-LD can be devolved easily from JSON-based systems.
- **Usable as RDF:** JSON-LD can be mapped to / from RDF and can be used as RDF without having any knowledge about RDF.

From the above, terseness and simplicity are the main reasons JSON-LD was chosen over RDF/XML. JSON-LD also allows for referencing external files to provide context. This means contextual information can be requested on-demand and makes JSON-LD better suited to situations with high response times or low bandwidth usage requirements. We think that using JSON-LD will reduce the complexity of VITAL development by (1) making it possible to reuse a large number of existing tools and (2) reduce the inherent complexity of RDF documents. Ultimately, this will increase VITAL's uptake and success. In the following we give a short overview of the main JSON-LD features and concepts. More information can be found in [SLK+14].

The data model underlying JSON-LD is a labelled, directed graph. There are a few important keywords, such as `@context`, `@id`, `@value`, and `@type`. These keywords are the core part of JSON-LD. Four basic concepts should be considered:

- **Context:** A context in JSON-LD allows using shortcut terms to make the JSON-LD file shorter and easier to read (as well as increasing its resemblance with pure JSON). The context maps terms to IRIs. A context can also be externalised and reused for multiple JSON-LD files by referencing its URI.
- **IRIs:** Internationalised Resource Identifiers (IRIs) are used to identify nodes and properties in Linked Data. In JSON-LD two kinds of IRIs are used: absolute IRIs and relative IRIs. JSON-LD also allows defining a common prefix for relative IRIs using the keyword `@vocab`.
- **Node Identifiers:** Node identifiers (using the keyword `@id`) reference nodes externally. As a result of using `@id`, any RDF triples produced for this node would use the given IRI as their subject. If an application follows this IRI it should be able to find some more information about the node. If no node identifier is specified, the RDF mapping will use blank nodes.
- **Specifying the Type:** It is possible to specify the type of a distinct node with the keyword `@type`. When mapping to RDF, this creates a new triple with the node as the subject, a property `rdf:type` and the given type as the object (given as an IRI).

2.3 SPARQL

Assuming that there is RDF data, then a developer needs a language to query it. SPARQL (SPARQL Protocol and RDF Query Language) is able to query single RDF graphs and it is used to retrieve and manipulate data, which is stored in RDF. There are four query variations that SPARQL can distinguish: `SELECT`, `CONSTRUCT`, `ASK` and `DESCRIBE` queries. The most basic construct of a SPARQL query are graph patterns, explained in [HaSe13]. A basic graph pattern is similar to a RDF triple except that the subject and predicate can be variables as well. The example in Figure 1 shows a query returning the names of all pairs of people that know each other.

```
PREFIX foaf: <http://xmlns.com/foaf/0.1/>

SELECT ?name1 ?name2
FROM <http://example.org/foaf>
WHERE {
    ?person1 foaf:knows ?person2.
    ?person1 foaf:name ?name1.
    ?person2 foaf:name ?name2
}
```

Figure 1: SPARQL Example [LWS13]

2.4 Ontologies

In addition to RDF and SPARQL another very important technology in Linked Data are ontologies. This chapter is a brief introduction into ontologies and their meaning.

Linked Data can be used in cases where data originates from different sources. To integrate all data, be it from one or different sources, there have to be some rules. Some rules determine how the RDF graph is to be built and how triples may be connected or not. These rules are given by ontologies. An ontology specifies formally the conceptualisation of a domain of interest. As the conceptualisation is formal, a computer can automatically reason on it. There are practical ontologies for different domains of interest. An ontology consists of concepts (also referred to as classes), relations (also called properties), instances and axioms. It defines basic terms and relationships.

To specify ontologies, the W3C published the Web Ontology Language (OWL²), which builds on RDF. OWL facilitates mechanism for creating concepts, instances, relations and axioms. Concepts can have super and sub concepts. Axioms provide information about classes and properties. This topic is explained in detail by [DSW06] and [NoMc01].

Reuse of ontologies is crucial. If some new data models arise they should be attached to an existing ontology. This ontology grows by doing so and helps users on any of its nodes to reach every other node in this ontology graph. The user will get much more information as just of his own model – if he/she wants to.

2.5 Conclusion

Linked Data (using RDF, SPARQL and ontologies) helps to describe and integrate data that is provided by different organisations in an interoperable way. This is ideally suited for VITAL. VITAL envisages that a multitude of (independent) organisations and entities deploy and operate different sensors and IoT systems, which produce data (and offer functionality) that should be integrated in a platform agnostic way.

² www.w3.org/2004/OWL

VITAL uses Linked Data standards for modelling and accessing data. RDF is used as the basic data model. JSON-LD is used as the data format. SPARQL is used to query data using complex sub graph patterns. Finally, ontologies are used to specify our data formally. To use the potential of Linked Data for interconnecting with as many external information sources as possible (making their information readily available for VITAL developers), we use the good practice of reusing as many ontologies as possible. We require ontologies in different areas. First, VITAL must allow to model sensors and sensor measurements, which are the basis of any IoT system. Second, VITAL has to provide means to model entities that are relevant to Smart Cities. Third, VITAL must allow to model IoT systems and services that are integrated into the VITAL platform. And finally, it must provide ontologies to model the VITAL system itself. In the remainder of this deliverable we discuss how to model these different parts using which ontologies. Table 1 provides an overview of the main ontologies used in VITAL, their namespaces and the prefixes that we use to refer to them.

Prefix	Ontology / Language	Namespace
dcn	Delivery Context ontology	http://www.w3.org/2007/uwa/context/deliveryContext.owl#
dul	DOLCE+DnS Ultralite ontology	http://www.ontologydesignpatterns.org/ont/dul/DUL.owl#
geo	Basic Geo (WGS84) ontology	http://www.w3.org/2003/01/geo/wgs84_pos#
hrest	hRESTS ontology	http://www.wsmo.org/ns/hrests#
msm	Minimal Service Model ontology	
owl	Web Ontology Language	http://www.w3.org/2002/07/owl#
rdfs	RDF Schema ontology	http://www.w3.org/2000/01/rdf-schema#
sawSDL	Semantic Annotations for WSDL and XML Schema ontology	http://www.w3.org/ns/sawSDL#
ssn	Semantic Sensor Network ontology	http://purl.oclc.org/NET/ssnx/ssn#
time	OWL Time ontology	http://www.w3.org/2006/time#
vital	VITAL ontology	http://vital-iot.com/ontology#
wsl	WSMO-Lite ontology	http://www.wsmo.org/ns/wsmo-lite#
xsd	XML Schema Definition	http://www.w3.org/2001/XMLSchema#
qudt	Quantities, Units, Dimensions and Data Types Ontologies	http://qudt.org/schema/qudt#
foaf	Friend of a Friend	http://xmlns.com/foaf/
s4ac	Social Semantic SPARQL Security for Access Control	http://ns.inria.fr/s4ac/v2#
otn	Ontology of Transportation Networks	http://www.pms.ifi.lmu.de/reverse-wga1/otn/OTN.owl

Table 1: Ontology/ Language Prefixes

3 SENSORS AND MEASUREMENTS

Networked sensors and their measurements are one of the most important parts of any data model for IoT systems and therefore also an integral part of the VITAL data model. To better define the semantics of IoT data, a number of ontologies have been developed on multiple layers of abstraction [w3c]:

- (1) **sensor-centric ontologies** like the Ontonym sensor ontology [SKD+09], the Sensor Data Ontology [ELS07], and OntoSensor [GoRu06];
- (2) **observation-centric ontologies** like the Semantic Sensor Network (SSN) Ontology [CBB+12], the Sensei Observation and Measurement Ontology [WeBa09], and stimuli-centered ontologies [SJB+09], as well as
- (3) **context-centric ontologies** like COMANTO [SRA06] and SOUPA [CPF+04].

Clearly it is impossible to specify a single ontology that defines the semantics of all possible data items, as they are in many cases application (domain) specific. This has led to the development of rather abstract and complex ontologies that try to fit all possible cases by providing a conceptual framework only, omitting concrete instances like specific sensor models, etc. Such ontologies try to impose an overarching structure onto IoT systems and their data, e.g. specifying abstract metadata classes for stimuli, observations, measurements, sensors and features of interest.

In practice such ontologies must be combined with additional ontologies to define concrete instances of abstract concepts. As an example, while a generic sensor ontology may specify how to model what a sensor is measuring, additional definitions must be used to model a concrete location sensor.

In the following we first present, SSN as the generic sensor ontology that the VITAL consortium has selected to be used in VITAL. We then discuss a number of further ontologies as well as additional definitions that are used in VITAL to create concrete data items.

3.1 The W3C Semantic Sensor Network Ontology

The Semantic Sensor Network (SSN) ontology [CBB+12] defines a conceptual framework for describing sensors and observations. It was developed by the W3C Semantic Sensor Network Incubator group (SSN-XG). In the following the ontology will be explained more in detail to give an understanding about its meaning by presenting what it describes, the used pattern (SSO) and the four main perspectives.

Overall the SSN ontology is able to describe:

- sensors including their accuracy and capabilities,
- observations,
- methods for sensing,
- concepts for operating and survival ranges, and
- deployments.

A sensor could be anything that observes, be it an electronic object, a virtual object or a human. The ranges are used in the definition of sensors conjoined with the performance of these sensors. The description of deployment includes the deployment lifetime as well as the sensing purpose of the deployed macro instrument.

3.1.1 Stimulus-Sensor-Observation Pattern

The ontology is built around an ontology design pattern called Stimulus-Sensor-Observation (SSO) Pattern. This pattern describes the relationships between sensors, stimuli and observations.

A sensor is modelled as a physical object (`dul:PhysicalObject`) and is able to observe, transform and represent data. How it observes is defined in sensing methods (imported from other ontologies).

A stimulus (`dul:Event`) is represented by a change or state which can be detected and used by a sensor to measure some property. It is comparable to a proxy for an observation property. A property, in turn, is an observable characteristic of a real-world entity (`ssn:FeatureOfInterest`).

Observations are the connectors in the SSO pattern. They can link between the act of sensing, a stimulus event, a sensor, a method, a result, an observed feature and a property – and put them in an interpretative context. An observation is modelled as a social construct (from `dul:Situation`).

3.1.2 Perspectives

For a better understanding of the ontology in terms of sensors and observations, there are four main perspectives in the SSN ontology, explained in the following.

Sensor perspective

A sensor is described with a stimulus, a sensing method, observations and capabilities. The performance of a sensor can be influenced by the environment. This is referred to as measurement capabilities. Such capabilities are i.e. accuracy, measurement range, measurement precision, and measurement resolution. One sensor can have many measurement capabilities, representing the capabilities of the sensor in different conditions (observable properties of the sensor-environment).

Observation perspective

The observation perspective describes an observation. An observation includes a context for interpreting incoming stimuli and puts the observation event (including the stimulus) into an interpreting context. A context includes observed features, properties, the observing sensor, the result and the used sensing method. As an observation is a social construct, a stimulus event can be abstracted from its (potential) interpretations. The sensing method can be a principle underlying a sensor or a description of how the observations were done.

System perspective

A system can have sub-systems or sub-concepts like devices and sensing devices. There are operating and survival restrictions for a system, similar to the ones for sensors and measurement capabilities, which are observable properties of the system. The operating range specifies the range in which the system aims to operate. The survival range is a range in which a sensor can bare without release a lasting damage. The process of combining all life-cycle phases of a deployed system is named deployment. It includes installation, maintenance and decommissioning of a system.

Feature and property perspective

The feature and property perspective focuses on properties, more specifically on the sensors that sense some distinct property or on the observations that were made about a distinct property.

3.1.3 Conclusion

To sum up, the SSN ontology was built to describe sensors, sensing and measurement capabilities of sensors as well as the resulting observations and deployment. The ontology covers big parts of the SensorML and O&M standards. SSN restricts itself to a conceptual view. It does not include specific definitions for concrete sensors or domain areas. Therefore, to realise a concrete IoT system based on the SSN ontology additional domain specific ontologies must be imported. For example you would use a temperature ontology on top of the SSN ontology to model a temperature sensor. VITAL uses the SSN ontology as the basis for modelling all entities and activities related to sensor data and metadata in a Smart City. Necessary extensions with additional ontologies (e.g. for location modelling or domain specific measurements) are discussed in the remainder of this chapter.

3.2 Basic Concepts and Ontologies

In this section we present ontologies to model basic concepts like time, location and unit of measurement. These will be used throughout the VITAL system. After that we discuss how to model sensors and sensor measurements in more detail.

3.2.1 Time

Temporal aspects are essential for any system addressing real world phenomena, e.g. smart city IoT systems. Timestamps can be used to describe when a sensor reading was taken or when it was valid. Multiple readings can be ordered by the time of their occurrence. Clients may specify that they are interested in the current state of an environment or in the state it had one week ago. To model this, VITAL has to provide an ontology for time as well as temporal properties and relations. A well-established ontology for this is OWL Time [HoPa06]. OWL Time allows describing of temporal properties and relationships. It also supports time intervals as well as durations, which are e.g useful when describing imprecise measurement times as well as complex event specifications.

In VITAL, all timestamps, temporal properties and relations are described using OWL Time. An example for a timestamp in a sensor measurement is given in Figure 2.

```

{
  "@context": {
    "time": "http://www.w3.org/2006/time#",
    "xsd": "http://www.w3.org/2001/XMLSchema#",
    "time:inXSDDateTime": {
      "@type": "xsd:dateTime"
    }
  },
  "@type": "time:Instant",
  "time:inXSDDateTime": "2014-08-20T16:47:32+01:00"
}

```

Figure 2: Example Time Instant Specification with OWL Time

3.2.2 Location

Locations in the physical world are another basic concept that has to be modelled in VITAL. There is a multitude of different location models and ontologies available today, including geographical and symbolic location models. VITAL follows a practical approach to allow easy usage of the system while at the same time being flexible enough for advanced use cases.

WGS84 coordinates are used as the basic location model, since they are the de-facto standard for outdoor localisation using the GPS system. To model them, the basic Geo (WGS84) vocabulary³ is used. See Figure 3 for an example. In addition, symbolic names are often used as locations. VITAL allows using symbolic names instead of WGS84 coordinates. VITAL uses the Linked GeoData system⁴ to model more complex location concepts, including symbolic names, cell-based locations and inter-location relationships. To map between symbolic names and WGS84 coordinates, VITAL uses the GeoNames system⁵. If a location is given as a symbolic name, VITAL automatically tries to retrieve the WGS84 coordinates for it and enriches the location data with them. The same is true the other way round.

```

{
  "@context": {
    "geo": "http://www.w3.org/2003/01/geo/wgs84_pos#"
  },
  "@type": "geo:Point",
  "geo:lat": "55.701",
  "geo:long": "12.552",
  "geo:alt": "4.33"
}

```

Figure 3: Example WGS84 Location with the Basic Geo Vocabulary

³ <http://www.w3.org/2003/01/geo/>

⁴ <http://linkedgeodata.org/>

⁵ <http://www.geonames.org/>

3.2.3 Unit Of Measurement

Different properties in the VITAL data models represent physical magnitudes like “length” or “weight”. Each one of these properties should be associated with an unambiguous unit of measurement, e.g. “metre” or “kg”. Otherwise, cultural differences may lead to clients interpreting values incorrectly, e.g. by assuming a length is given in “feet”, while it is actually given in “metre”. While RDF allows specifying the type of a given property value by tagging it with a special type identifier (e.g. “13”^{^^}<<http://www.w3.org/2001/XMLSchema#int>>), it does not support tagging values directly with a unit of measurement. At the same time, there is currently no single accepted ontology to model units of measurements in linked data. A number of potential ontologies were found and four were chosen for detailed evaluation. We discuss these four in the following.

Quantities, Units, Dimensions and Data Types Ontologies (QUDT) [HKHS14] is a family of ontologies developed by TopQuadrant and sponsored by NASA designed to formalize quantities, units of measurement, dimensions and types in RDF/OWL formats. Due to this ambitious goal it is incredibly broad and precise. It models base types such as length and time and builds derived types as a hierarchy (e.g. $\text{Velocity} = \text{Length} / \text{Time}$, $\text{Kilometres Per Hour} = \text{Kilometres Travelled} / \text{Time taken}$). Where appropriate, it also references similar ontologies with `sameAs` and `exactMatch` relationships as well as DBpedia entries where appropriate. This gives items associated with QUDT types a huge amount of semantic information. Similar to QUDT, the **Units of Measurement Ontology (UO)** [UO] establishes a hierarchy of base and derived types. It contains a large number of types but is missing some common derived types such as “Kilometres Per Hour”. It also fails to link to external resources such as DBpedia where relevant which makes it less useful than QUDT. The **Ontology of units of Measure (OM)** [OM] stands out from the other evaluated ontologies by using URIs as the value for the majority of each type’s information. This immediately gives feedback to the user that values are unique. Thanks to each type using the same signature, this lets the user easily compare types for similarities using simple methods such as string comparison rather than needing to traverse the type hierarchy. Each of these URIs (e.g. `om:kilometre`, `om:dimension`, `om:length-dimension`) can be followed to obtain more information including other users of this value. OM has a large number of derived types available but – just as UO – lacks links to external resources such as DBpedia and other ontologies. Finally, the **Measurements Unit Ontology (MUO)** [PoBe08] aims to solve the same problem as the other evaluated ontologies by establishing a hierarchy of base and derived types. However, it appears to have been abandoned before completion as of time of writing.

VITAL chooses QUDT as the ontology for units of measurements due to the impressive scope and amount of information available on each type as well as the reputation of the maintainer and sponsoring party. QUDT is also actively maintained, with the newest version being released in March 2014. In addition, the links to DBpedia provide a tight integration with a lot of further data items, both DBpedia itself and in data sources linking to it – even more so since VITAL will use DBpedia to link to information about cities (see Section 4.1). However, due to the comparable strength of OM, a re-evaluation is recommended in the future to ensure QUDT remains the best choice. An example for using QUDT is given in Figure 6.

3.3 Sensors

After presenting ontologies for basic concepts we can now discuss how to model sensors, sensor measurements and their descriptions in VITAL. In general, VITAL reuses and extends the SSN ontology and the Delivery Context (DC) ontologies.

A sensor is modelled as a `VitalSensor`, a subclass of `ssn:Sensor` and `dcn:Device`⁶. According to the SSN, a “sensor can do (implements) sensing: that is, a sensor is any entity that can follow a sensing method and thus observe some Property of a FeatureOfInterest. Sensors may be physical devices, computational methods, a laboratory setup with a person following a method, or any other thing that can follow a Sensing Method to observe a Property” [ssn]. By using the SSN ontology, VITAL can immediately describe sensors in detail, including aspects like the properties that they observe, sensor locations, and sensor observations. The SSN ontology also allows to model non-functional aspects of a sensor, e.g. its accuracy or reliability, by adding a `ssn:hasMeasurementProperty` property to the sensor description that points to a `ssn:MeasurementCapability` (or a subclass of it). The reader is referred to the SSN ontology specification at [ssn] for more information about the ontology. The DC ontology defines a `dcn:Device` as a class that “represents a device in the delivery context” [FoLe09]. By using the DC ontology, VITAL can reuse a highly detailed set of ontologies describing many aspects of devices, including their software, their hardware as well as their networking.

Class: IotSystem

SubClassOf: `ssn:Sensor`, `dcn:Device`

3.3.1 Additional Ontology Definitions

In addition to the SSN and DCO ontologies, VITAL defines an additional property for sensors, `hasLastKnownLocation`. This property is a sub property of `dul:lastLocation` as specified in the SSN ontology description. It links to a location, which is the last known location of the sensor. Note that this property replaces `dul:hasLocation` to make the property’s semantic less ambiguous in VITAL. The property does not imply that this is the actual *current* location of the sensor. If the sensor is mobile, it could have moved to a new location after the description was created. If the property is not available in a sensor description, then the location of the sensor may not be known. However, a client may be able to access the current location using a so-called localizer service (see Section 6.5.4).

ObjectProperty: hasLastKnownLocation

SubPropertyOf: `dul:hasLocation`

Domain: `VitalSensor`

Range: `dul:Entity`

⁶ Note that `dcn` is the prefix of one of the DC ontologies.

Note that the location of a sensor can be modelled with different types as specified before, e.g. as a `geo:Point` in case GPS coordinates are used. To be as flexible as possible, we use the generic `dul:Entity` class to represent all different location types here. It is taken from the DOLCE+DnS Ultralite (DUL) ontology and defines it as “[a]nything: real, possible, or imaginary, which some modeller wants to talk about for some purpose.” [dul]

In addition to the basic description of a sensor, some VITAL services may require additional information. Properties and classes to model this information are described in Chapter 6 of this deliverable.

3.3.2 JSON-LD Definitions

Note that since Linked Data in VITAL is always formatted as JSON-LD we can introduce some additional definitions (in a JSON-LD `context` section as shown in Figure 4) that do not change the used ontology or the resulting RDF triples but align the JSON-LD representation more closely to ‘normal’ JSON and thus makes it easier for developers to work with the data. We use this approach repeatedly in this document. As can be seen in the figure, we first specify that all JSON keys that do not specify a prefix will be expanded to URIs in the VITAL ontology namespace. This results in more compact files with less clutter. Then we define that the key `uri` will be mapped to a JSON-LD node identifier (`@id`). The node identifier is used to create the URI that is used as the subject in RDF triples.

```
{
  "@context": {
    "@vocab": "http://vital-iot.com/ontology#",
    "uri": "@id",
    "type": "@type",
    "name": "rdfs:label",
    "description": "rdfs:comment",
    "vital": "http://vital-iot.com/ontology#",
    "lsm": "http://lsm.derii.ie/ont/lsm.owl#",
    "ssn": "http://purl.oclc.org/NET/ssnx/ssn#",
    "geo": "http://www.w3.org/2003/01/geo/wgs84_pos#",
    "msm": "http://iserve.kmi.open.ac.uk/ns/msm#",
    "hrest": "http://www.wsmo.org/ns/hrests#",
    "hard": "http://www.w3.org/2007/uwa/context/hardware.owl#",
    "soft": "http://www.w3.org/2007/uwa/context/software.owl#",
    "hrest:hasAddress": {
      "@type": "hrest:URITemplate"
    },
    "ssn:madeObservation": {
      "@type": "@id"
    }
  }
}
```

Figure 4: JSON-LD Context for Sensor

Similarly, we specify that the key `type` will be mapped to the JSON-LD keyword `@type`. This results in an RDF triple being created that specifies the RDF type of a node. To further simplify the JSON-LD file, we specify that the key `name` will be

mapped to `rdfs:label` and the key description will be mapped to `rdfs:comment`. We then specify a number of prefixes that can be used in the JSON-LD description to reduce the length of keys by specifying them as so-called terms. Terms are automatically expanded using the provided prefix URI. As an example, the term `geo:lat` would be expanded to `http://www.w3.org/2003/01/geo/wgs84_pos#lat`. Finally we specify type specifications that declare that a value given for the key `ssn:madeObservation` or `hrest:hasAddress`⁷ should be mapped to a node id instead of a string. All these mappings are completely transparent to developers and can be ignored by clients. They are only relevant if the JSON-LD file is mapped to RDF triples internally. Together, they reduce the complexity of the resulting JSON-LD file and make it both smaller and easier to read and understand for JSON developers.

3.3.3 Sensor Description Example

To illustrate the resulting sensor descriptions, we give a short example for a sensor description (in JSON-LD format) in Figure 5. The example uses an external version of the context specification for VITAL sensors that we have just described. This further reduces the size of the JSON-LD file. VITAL clients can access the external context file at `http://vital-iot.org/contexts/sensor.jsonld` if necessary. However, we expect that this will not be necessary for most scenarios and normal operation. The context file can also be cached on the client side, because it is independent of a specific VITAL sensor and can be reused.

```
{
  "@context": "http://vital-iot.org/contexts/sensor.jsonld",
  "name": "TemperatureSensor No.123",
  "type": "VitalSensor",
  "description": "This is an example sensor",
  "uri": "http://www.example.com/vital/sensor/123",
  "hasLastKnownLocation": {
    "type": "geo:Point",
    "geo:lat": "53.2719",
    "geo:long": "-9.0489"
  },
  "ssn:observes": [
    {
      "type": "http://lsm.derri.ie/OpenIoT/Light",
      "uri": "http://www.example.com/vital/sensor/123/light"
    },
    {
      "type": "http://lsm.derri.ie/OpenIoT/Temperature",
      "uri": "http://www.example.com/vital/sensor/123/temperature"
    }
  ],
  "ssn:madeObservation":
    "http://www.example.com/vital/sensor/123/obsvn/1"
}
```

Figure 5: Example Sensor Description

⁷ see Section 5.2 for more information about *hrest* and the hRESTS ontology.

In the example, we describe a sensor with its name, its type (a `VitalSensor`), its URI in the VITAL system and its description. The sensor provides a last known location as a GPS (WGS84) location (of type `geo:Point`). It also specifies that it observes two different parameters, the light level and the temperature, as well as URIs to retrieve current measurements for both. Finally, a sensor *may* specify observations it made in the past, in this case an observation with ID `http://www.example.com/vital/sensor/123/obsvn/1`.

3.4 Sensor Measurements

Similarly to sensors, VITAL uses the SSN ontology to model sensor measurements. A measurement is modelled as an `ssn:Observation`. An example can be seen in Figure 6. Figure 7 shows the JSON-LD context used by the example. The observation contains a link to an observed property (using `ssn:observationProperty`) to specify what the observation is measuring. In addition, it specifies when the measurement was taken (`ssn:observationResultTime`), at which location (`dul:hasLocation`) in WGS84 format, the quality of the measurement (`ssn:observationQuality`), as well as the measured value (`ssn:observationResult`), 21.0 Degree Celcius, with the unit of measurement specified with the QUDT ontology.

```
{
  "@context": "http://vital-iot.org/contexts/measurement.jsonld",
  "uri": "http://www.example.com/vital/sensor/123/obsvn/1",
  "type": "ssn:Observation",
  "ssn:observationProperty": {
    "type": "http://lsm.derivi.ie/OpenIoT/Temperature"
  },
  "ssn:observationResultTime": {
    "inXSDDateTime": "2014-08-20T16:47:32+01:00"
  },
  "dul:hasLocation": {
    "type": "geo:Point",
    "geo:lat": "55.701",
    "geo:long": "12.552",
    "geo:alt": "4.33"
  },
  "ssn:observationQuality": {
    "ssn:hasMeasurementProperty": {
      "type": "Reliability",
      "hasValue": "HighReliability"
    }
  },
  "ssn:observationResult": {
    "type": "ssn:SensorOutput",
    "ssn:hasValue": {
      "type": "ssn:ObservationValue",
      "value": "21.0",
      "qudt:unit": "qudt:DegreeCelsius"
    }
  }
}
```

Figure 6: Example Measurement

```
{
  "@context": {
    "@vocab": "http://vital-iot.com/ontology#",
    "uri": "@id",
    "type": "@type",
    "name": "rdfs:label",
    "description": "rdfs:comment",
    "vital": "http://vital-iot.com/ontology#",
    "lsm": "http://lsm.derii.ie/ont/lsm.owl#",
    "ssn": "http://purl.oclc.org/NET/ssnx/ssn#",
    "geo": "http://www.w3.org/2003/01/geo/wgs84_pos#",
    "time": "http://www.w3.org/2006/time#",
    "qudt": "http://qudt.org/vocab/unit#",
    "xsd": "http://www.w3.org/2001/XMLSchema#",
    "time:inXSDDateTime": {
      "@type": "xsd:dateTime"
    },
    "qudt:unit": {
      "@type": "@id"
    }
  }
}
```

Figure 7: JSON-LD Context for Measurement

3.5 Conclusion

Sensors and sensor measurements are a central part of each IoT data model. VITAL uses (and extends) the well-established SSN ontology to model these information items, including meta-information like the type of measurements a sensor can perform or the quality of a measurement. To model basic concepts like time, location and unit of measurement, VITAL relies on additional ontologies, like OWL Time and the basic geo ontology. Together with the SSN ontology this provides the means to specify sensors and measurements in a platform agnostic way, as required by VITAL. It also allows us to link them with other systems and information sources that use the same or similar ontologies, like the FP7 project OpenIoT. This also enables the reuse of existing tools and software components for VITAL, without the need to implement everything from scratch. An example for this is the VITAL data management component, which will be based on OpenIoT.

However, we cannot yet describe real physical instances of sensors (or measurements), since we do not have a common ontology for specific values for our descriptions. As an example, while we know how to specify what property of the real world a sensor is observing (by using `ssn:observes`) we do not have a specification for modelling real instances of such properties, e.g. temperature, speed or light level. These are application domain specific. This issue is discussed further in the next chapter.

4 SMART CITIES

‘Smart City’ is quite a broad term that represents many different aspects of a city’s population, economy and infrastructure [CDN09]. For this document, we define a Smart City as one that uses ICT to integrate its different capabilities and enhance the quality of life for its population. Smart cities have many capabilities that must be modelled as part of a smart city ontology. The most active and relevant producers of data come under the following headings (see also VITAL deliverable D2.1 [SoKa13]):

- **Transport** – e.g. dynamic route calculation informed of accidents and congestion, integration of transport methods such as bus, tram and subway for public transport users.
- **Energy** – e.g. optimization of light levels, reporting of faults.
- **Emergency Services** – e.g. using real-time traffic data to inform emergency response units and to provide real-time information on accident and crimes.
- **Waste Management** – e.g. detection of full bins as well as automated reporting of missing/damaged bins.
- **Air and Water** – e.g. reporting of high pollution in air and water, radiation levels reporting of water leaks, automatic irrigation of parks and green areas.
- **Recreation** – e.g. provide data on local social events as well as nearby facilities, produce data on large events such as concerts and sports games to inform public transport and law enforcement systems of potential increased activity.

Other capabilities such as smart homes (Including advanced integration cases such as those designed for OAPs and the disabled) and smart business and industry are mostly consumers but may also provide alerts (e.g. Factories may provide notice when beginning a water-intensive task).

The VITAL use cases (see VITAL deliverable D2.2 [SoKa14]) so far focus on (smart) transport and traffic management. Therefore, in the following we discuss how to model data items and properties that are relevant for smart traffic scenarios. Clearly, VITAL is not restricted to smart transport scenarios. A user who would like to use VITAL for other smart city aspects can do so by specifying additional ontology elements. Due to the nature of Linked Data, these additional elements can be added at any time without the need to redesign the system. In the following we first discuss briefly how VITAL includes general information about cities that may be useful for different aspects of smart city applications. After that we focus on modelling smart transport related information.

4.1 Cities

VITAL obtains the majority of its semantic information on cities from DBpedia, using the classic DBpedia dataset for most information with the option of using DBpedia live for information that updates more frequently. It is also encouraged to link real places and services in cities back to DBpedia to improve the amount of knowledge available. For example, while Camden Road would be modelled as an `otn:Road` as part of a smart transport system, it should also link to `dbpedia:Camden_Road`.

4.2 Smart Transport

Smart Transport is an important topic in Smart Cities and covers a wide range of domains. Some examples include tracking pedestrian congestion in an area through footfalls per minute, smart traffic light systems or route recommendation to empty parking spaces capable of reserving a spot en route. There are also opportunities for integrated solutions with other smart city services such as automatic contact of emergency services in the case of an accident and the ability to detour traffic away from the scene through the use of traffic light control.

In this section, we'll first discuss how traffic infrastructure is modelled in VITAL before going on to discuss what kind of services can be built on top of this infrastructure. Finally, we'll discuss how VITAL will showcase the power of integrated IoT smart city solutions.

4.2.1 Transport Infrastructure

VITAL models transport infrastructure using a combination of ontologies. The core of these is the *Ontology of Transportation Networks* (OTN) [OTN]. This ontology allows easy modelling of a transport network graph with connections between infrastructure such as bus and train stations as well as events such as accidents and blocked passages.

Modelling public transport was one of OTN's main goals during its design so it is trivial to integrate different public transport methods. For example, a journey that requires a 20 min train journey, 3 minutes walking to a regional bus station, a 1 hour bus journey, 2 minutes walking to a city bus service, a 5 minute bus journey and 4 minutes walking to destination can be described as a route through a single graph. This graph representation of public transport can be used by services such as those described in Section 4.2.2 to plan journeys. Another example would be providing automatic detours in the case of accidents and construction and using an ambulance's projected route to change traffic lights to red based on its current location with dynamic recalculation if the ambulance changes from the suggested route.

VITAL also aims to extend OTN to allow new public and private transport types. Some common examples include locations of taxi offices and bicycle racks but this will also allow the addition of less common methods of transport such as rickshaws, which are extremely popular in some countries and tour buses, which may be preferable to tourists.

Each node in the OTN graph is also augmented with approximate longitude and latitude, corresponding DBpedia information and URIs for any sensors located at this location.

4.2.2 Transport Services

Route recommendation services aim to provide a user with an optimal route to their destination, generally using live traffic data and other available information such as construction works and accidents to reach this solution. There are a number of high-profile examples of these such as Google Maps, Bing Maps and Nokia's Here Maps. However, while these solutions attempt to take external factors such as those

previously mentioned into account they cannot take advantage of any knowledge of other drivers' routes. Smart cities do not need to suffer this limitation.

The Ontology based Traffic Jam Control System (OnTraJaC) [HaMa12] is designed to recommend routes to avoid traffic jams and support shorter trip times. It achieves this by taking into account all subscribed vehicles during route recommendation and is able to recommend a route to each user that does not conflict with others.

This approach should produce the optimal route for each subscribed user but also has some interesting opportunities for future expansion and integration with other Smart City services. For example, this system could be used to tweak subscribed user's routes when they conflict with the path of an emergency vehicle. In Section 4.2.1 it was suggested that traffic light infrastructure could be integrated with emergency response systems. However, this would not affect users driving along the emergency vehicles path, only those on other routes that would cross it at a junction. By combining this approach with the route recommendation service, subscribed users could be recommended to turn off the emergency vehicles projected route just before the emergency vehicle reaches their position. This would prevent the emergency vehicle having to potentially overtake them and would be safer for both subscribed users and emergency response teams.

While there are ontologies that support similar functionality to existing personal route recommendation systems – such as the Urban Traffic Ontology [UTO] – these should in theory suggest the same route as OnTraJaCS would for the case of $n = 1$, where n equals the number of subscribed users taken into account during route calculation. Therefore, we have decided there is no need for any other route recommendation ontology and instead we should simply accommodate a scenario where OnTraJaCS operates with knowledge of only one user.

To support public transport route planning many ontologies were examined but most have their own model of how the transport network should be modelled. Since VITAL has chosen to use a combination of ontologies to represent transport infrastructure and public transport facilities, chiefly OTN, it is non-trivial to integrate these ontologies with existing ones. It is therefore the recommendation of the VITAL consortium to wait until the model for a city's transport infrastructure and facilities is finalized before attempting to adapt existing ontologies such as those described in [HKO+10] and [CoWi01] to utilize VITAL's data model.

4.3 Conclusion

Modelling the infrastructure and available facilities and services of a smart city is a non-trivial task. VITAL therefore aims to cover one of the most important components: Smart Transport. This will include modelling the cities transport infrastructure and services such as route recommendation, but will also explore integration with other smart services such as emergency response to showcase the power of integrated smart city solutions. Future services can then be integrated with VITAL's traffic infrastructure from the beginning, which will allow scenarios such as automatic route calculation to entities such as burst water pipes and fallen electricity lines for repair teams. In subsequent versions of the VITAL data models, we intend to elaborate on more specific attributes and relationships that will serve as a basis for supporting the VITAL transport scenarios and use cases.

5 IOT SYSTEMS AND SERVICES

VITAL integrates existing IoT systems (i.e. deployed platforms) and allows clients (applications as well as VITAL system services) to access (meta-) data and services of such systems. So far the consortium has selected four platforms for which example deployments will be integrated as a proof of concept: GSN, Reply H1, INRIA FIT and Xively. To integrate systems and work with them, VITAL needs a set of models to describe IoT services, their data and services. For data, we will use the models presented in Chapter 3. In the remainder of this chapter we present the models used for an IoT system itself as well as for generic IoT services.

5.1 IoT Systems

VITAL models an IoT system as a subclass of `ssn:System` with a number of additional properties. An IoT system description always includes a basic set of properties that describe general aspects of the system, e.g. its operator. In addition, a system description may specify a set of *IoT services* that it offers.

Class: IotSystem

SubClassOf: `ssn:System`

5.1.1 General Metadata

To describe general metadata about the system, VITAL supports three new properties. These properties are:

- **status**, pointing to a status description for the system (this property must be available 0..1 times),
- **operator**, pointing to the entity that is responsible for operating this system (this property must be available 0..n times),
- **serviceArea**, pointing to the spatial context of this system, e.g. the city it is operating in (this property must be available 0..n times).

ObjectProperty: status

Domain: `IotSystem`

Range: `OperationalState`

ObjectProperty: operator

Domain: `IotSystem`

Range: `dul:Entity`

ObjectProperty: serviceArea

Domain: IotSystem

Range: dul:Entity

Note that `status` provides only coarse-grained aggregated information about the overall state of the system. More detailed information about the state of a system may be retrieved from an optional monitoring service. In case an IoT system does not provide any such service, the system description may need to be modified manually.

In addition, a `VitalSystem` supports properties that are independent of VITAL, namely `rdfs:label` (to specify a human readable name for a system) and `rdfs:comment` (to give a human readable description of a system).

Figure 8 shows an example for such an IoT system description. Note that the example is using the external JSON-LD context shown in Figure 9.

```
{
  "@context": "http://vital-iot.org/contexts/system.jsonld",
  "name": "Camden Reply System",
  "description": "This is a deployment of Reply H1 in Camden",
  "uri": "http://www.reply.com/camden",
  "operator": "http://www.ctu.com/",
  "serviceArea": "http://dbpedia.org/page/Camden_Town",
  "status": "vital:Running"
}
```

Figure 8: Example IoT System Description

In the example, the system description first associates a name (“Camden Reply System”) and a `description` (stating that this is a deployment of the Reply H1 IoT system for Camden Town) with this system. Note that – similar to before – `name`, `description` and `uri` are identifiers that have been specified in the context to make the JSON-LD file cleaner and easier to read for developers with previous experience with JSON. The context maps them to `rdfs:label`, `rdfs:comment`, and `@id` (JSON-LD node identifiers), respectively. In the example, the node identifier `http://www.reply.com/camden` identifies this system description and – following Linked Data good practice – specifies the location where a client can retrieve the description via HTTP.

The system in the example is operated by CTU (`operator`). The `serviceArea` is identified by its DBPedia URI. This allows a client to query the DBPedia for more information about Camden and to integrate this information. The system is currently in the Running state (see Section 5.1.2 below for a more detailed explanation of `OperationalState`).

```

{
  "@context": {
    "@vocab": "http://vital-iot.com/ontology#",
    "uri": "@id",
    "vital": "http://vital-iot.com/ontology#",
    "rdfs": "http://www.w3.org/2000/01/rdf-schema#",
    "msm": "http://iserve.kmi.open.ac.uk/ns/msm#",
    "hrest": "http://www.wsmo.org/ns/hrests#",
    "hrest:hasAddress": {
      "@type": "hrest:URITemplate"
    },
    "type": "@type",
    "status": {
      "@type": "@id"
    },
    "status": {
      "@type": "@id"
    },
    "hasOutput": {
      "@type": "@id"
    },
    "hasInput": {
      "@type": "@id"
    },
    "operator": {
      "@type": "@id"
    },
    "serviceArea": {
      "@type": "@id"
    },
    "name": "rdfs:label",
    "description": "rdfs:comment"
  }
}

```

Figure 9: JSON-LD Context Specification for Systems

5.1.2 OperationalState

OperationalState specifies the operational state of a system (or other entities, e.g. services). So far, the following states are defined in VITAL as sub classes of OperationalState:

- **Operational:** The entity is currently not running (and may remain in this state indefinitely) but could be started if needed (please note that this does not guarantee that starting it will succeed),
- **StartingUp:** The entity is in the process of starting and will (probably) be running soon,
- **Running:** The entity is currently operating and available to use,
- **ShuttingDown:** The entity is in the process of shutting down and will soon be stopped,

- **Unavailable:** The entity is currently not operating and cannot be started at this time. Usually this implies that a manual intervention is necessary to return the entity to the Operational state.

Should additional states be needed, they will be added in later versions of this deliverable.

Class: Operational

SubClassOf: OperationalState

Class: StartingUp

SubClassOf: OperationalState

Class: Running

SubClassOf: OperationalState

Class: ShuttingDown

SubClassOf: OperationalState

Class: Unavailable

SubClassOf: OperationalState

5.1.3 Provided IoT Services

In addition to the metadata discussed so far, an IoT system may offer a set of IoT services to access its functionalities. Typical examples of such services are data access services, data stream management, monitoring, discovery, etc. Each system may offer different services, making it impossible to define a fixed set of services. VITAL therefore provides the means to specify a large variety of IoT services with a generic service specification ontology. More information about this can be found in Section 5.2. To allow an IoT system to link to descriptions of provided IoT services, VITAL introduces a new property `providesService`.

ObjectProperty: providesService

Domain: IotSystem

Range: msm:Service

Figure 10 shows an example for an extended description of an IoT system that provides two IoT services.

```

{
  "@context": "http://vital-iot.org/contexts/system.jsonld",
  "name": "Camden Reply System",
  "description": "This is a deployment of Reply H1 in Camden",
  "uri": "http://www.reply.com/camden",
  "status": "vital:Running",
  "operator": "http://www.ctu.com/",
  "serviceArea": "http://dbpedia.org/page/Camden_Town",
  "providesService": [
    {
      "type": "StreamManager",
      "msm:hasOperation": [
        {
          "type": "RegisterStream",
          "hrest:hasAddress": "http://www.ctu.com/vital/streamMng/",
          "hrest:hasMethod": "hrest:POST"
        },
        {
          "type": "DeregisterStream",
          "hrest:hasAddress":
            "http://www.ctu.com/vital/stream/{id}/deregister/",
          "hrest:hasMethod": "hrest:GET"
        }
      ]
    },
    {
      "type": "IcoAccess",
      "msm:hasOperation": {
        "hrest:hasAddress": "http://www.ctu.com/vital/ico/{id}",
        "hrest:hasMethod": "hrest:GET"
      }
    }
  ]
}

```

Figure 10: IoT System Description with Provided Services

The first service (of type `StreamManager`) allows clients to register and deregister for continuous updates from linked data streams. The second service (of type `IcoAccess`) allows clients to retrieve the current state of ICOs.

In the following we describe IoT services and the ontologies used for them in more detail. To do so, we first discuss existing ontologies. Then, we specify how to describe IoT services in VITAL.

5.2 IoT Services

In VITAL, an IoT system does not only provide access to IoT data (e.g. sensor measurements) but may offer a set of distinct and heterogeneous IoT services. An IoT service may be generic, e.g. a service to discovery ICOs or to access filtered data, or application specific, e.g. a service to reserve a parking space in a Smart City IoT system. In fact, VITAL models all functionality that can be exposed by an IoT

system and can be accessed and used by a client as an IoT service, including data access, e.g. reading a sensor measurement. VITAL therefore specifies a flexible data model to specify all different kinds of IoT services. This allows extending the system in case a provider of an IoT system wants to expose new application specific services. In addition to this, VITAL provides concrete instances of IoT service specifications for generic (system) IoT services like the aforementioned discovery and filtering services. Service providers *can* extend these instances for their specific IoT system, but they *should* use instances specified by VITAL to model services if possible. In the following we briefly discuss different existing ontologies and modelling approaches for services, before describing the VITAL IoT service model in more detail.

5.2.1 Service Ontologies And Semantic Description Languages

There are a number of existing ontologies and specification languages to model services, among them the Web Service Definition Language (WSDL) and the Web API Definition Language (WADL) for specifying (SOAP-based) web services and (REST-based) web APIs (also known as RESTful services). Both WSDL and WADL are not using ontologies. Ontology-based approaches (see [VHM+14] for a recent survey) can be classified in three main groups: approaches that aim at extending generic WSDL-based service descriptions with semantic data, approaches that concentrate on extending descriptions of RESTful services with semantic data, and approaches that concentrate on providing a conceptual framework for service descriptions. In the following we briefly discuss some of the existing approaches and discuss their applicability to the VITAL use case. In general, VITAL aims at modelling IoT services with a focus on simplicity, minimalism, reuse as well as support from an active community.

WSDL-based Semantic Service Description Languages

WSDL-based approaches to describe semantic web services, like OWL-S [MBH+04], Semantically Annotated WSDL (SAWSDL) [FaLa07], and the Web Service Modelling Ontology (WSMO) [LPR05] focus on generic and complex semantic extensions of WSDL to describe SOAP-based web services. As an example, OWL-S consists of three parts that are used to specify three aspects of web services: the service profile (specifying what the service does for potential clients, e.g. name, description and quality of service levels), the process model (specifying how to use a service on a semantic level, e.g. input and output messages), and the service grounding (specifying details on how to access a service on a technical level, e.g. protocols and addresses). To fully specify the service grounding, OWL-S relies on an additional (external) specification. Most commonly WSDL is used. Despite the long availability of approaches such as OWL-S (having been published as a W3C member submission in 2004), they have received limited uptake and do not take into account newer developments and technologies such as REST.

Light-weight Semantic Service Description Languages

To accommodate for this, a number of newer specifications have been developed. Microformats such as Semantic Annotations for REST (SA-REST) [GRS10] and HTML for restful Services (hRESTS) [KGV08] embed additional metadata into (X)HTML-based textual descriptions of (usually RESTful) web services. This

metadata can be used to describe services in more detail and allows creation of machine understandable descriptions that can be used for discovery and automatic orchestration. MicroWSMO [KVF09] extends hRESTS with the ability to enhance a service description with external semantic descriptions as well as transformation routines to map between raw and semantic data. It is similar to SAWSDL but simplified to be used for RESTful services, only.

Higher Layer Conceptual Integration Frameworks

WSMO-Lite [VKVF08] is a lightweight subset of WSMO that provides a number of semantic annotations to describe web services. It can be used in conjunction with different lower layer ontologies and languages, e.g. on top of SAWSDL/WSDL for SOAP services and on top of MicroWSMO/hRESTS for RESTful services.

The Minimal Service Model (MSM)⁸ [KGV08] [PML+10] is a very lightweight ontology for the semantic modelling of Web service descriptions. It consists of only a few general concepts (e.g. a service and an operation) that are not detailed further in MSM. Instead, MSM provides a conceptual framework to integrate semantic descriptions that have been created using different other approaches. To actually specify a service, MSM must be combined with other ontologies, e.g. hRESTS, MicroWSMO and WSMO-Lite. While WSMO-Lite integrates semantic annotations, MSM integrates the basic concepts of a service description.

5.2.2 VITAL IoT Service Model

There is currently no single, standardised way to model IoT services. Based on the related work discussed before, the VITAL consortium decided to base its semantic IoT service model on existing work in the domain of web services. As discussed before, VITAL aims at providing a semantic model that is generic – yet simple and minimal, reuses existing ontologies as much as possible and allows to link with an active community as well as other current projects.

After careful consideration, the consortium selected to use the MSM as the basis of its IoT modelling ontology. MSM is small and easy to understand while at the same time providing integration with other languages like SAWSDL, WSMO-Lite and hRESTS. MSM is also widely used, e.g. by SOA4All [KNSP09], iServe [isrv], as well as the FP7 project COMPOSE [com]. COMPOSE focuses on transforming the IoT into an Internet of Services. By aligning VITAL with COMPOSE, applications using the VITAL platform will be able to participate in the open service marketplace envisioned in COMPOSE. At the same time, VITAL will be able to easily integrate all services offered by the COMPOSE marketplace. In addition, using MSM also allows a direct integration with iServe [isrv], an online service warehouse supporting service publication, analysis, and discovery. Nevertheless, the consortium will closely monitor all usage and experiences with the IoT service ontologies and may extend the used model to integrate further ontologies, e.g. OWL-S.

With this in mind, we can specify the following: In the VITAL system an IoT service is modelled as a RESTful (web) service that is described by Linked Data using the MSM ontologies. This allows publishing a description of the IoT service that can e.g. be used for discovery or for automatic composition tasks. An example for such an IoT

⁸ <http://iserve.kmi.open.ac.uk/ns/msm/msm-2013-05-03.html>

service is given in Figure 11 and Figure 12. Other than for earlier examples, to showcase how the JSON-LD file is mapped to RDF triples, we provide the same example twice, first in N3⁹ notation, then in JSON-LD format using the context shown in Figure 13.

```

@prefix msm: <http://iserve.kmi.open.ac.uk/ns/msm#>
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#>
@prefix hrest: <http://www.wsmo.org/ns/hrests#>
@prefix sawsdl: <http://www.w3.org/ns/sawsdl#>
@prefix wsl: <http://www.wsmo.org/ns/wsmo-lite#>
@prefix ex: <http://vital-iot.com/example#>
@prefix ibb: <http://www.ibb.gov.tr/ont/>

ex:parking a msm:Service;
rdfs:label "Parking Space Management Service";
  msm:hasOperation ex:getParkingStatus;
  sawsdl:modelReference ibb:ParkingSpaceService.

ex:getParkingStatus a msm:Operation;
  rdfs:label "getParkingSpaceStatus";
  hrest:hasAddress
    "http://www.ibb.gov.tr/parking/{id}"^^hrest:URITemplate;
  sawsdl:modelReference ibb:ParkingSpace;
  hrest:hasMethod hrest:GET;
  msm:hasOutput ibb:ParkingSpaceStatusReply.

ibb:ParkingSpaceStatusReply a wsl:Message;
  sawsdl:modelReference ibb:ParkingSpace.

```

Figure 11: Example IoT Service for Parking Space Management (N3)

In the example, we specify an IoT service for managing a parking space in Istanbul that allows a user to request the current status of parking spaces using their identifier (ID) by sending a HTTP GET request to the URI representing the parking space. Using MSM (red text), we first specify a service `ex:parking` that has an operation `ex:getParkingStatus`. In addition, we assign a label to the service that helps to describe the service for human readers. Using MicroWSMO (brown text) we also specify a reference to an external semantic description with the identifier `ibb:ParkingSpaceService` that describes the service in a machine understandable way. Note that the prefix used for MicroWSMO points to the

⁹ <http://www.w3.org/DesignIssues/Notation3.html>

SAWSDL ontology because MicroWSMO is reusing properties from it. The operation again specifies a label intended for human readers and in addition specifies how to access the operation using the hRESTs ontology (green text) by giving its address template and the used HTTP method (in this case GET). The address template allows to specify not only a single, fixed URI but a set of URIs (of type `hrest:URITemplate`) that can be used with this operation. In this example, the address template specifies that the operation can be called with the ID of parking spaces. As an example, calling the operation with address `http://www.ibb.gov.tr/parking/1` would get the status of the parking space with ID 1. The fact that the address models a parking space is again specified with MicroWSMO. The operation also has one reply that is specified as a message (using WSMO-Lite, violet text) and that is further specified semantically using MicroWSMO.

```
{
  "@context": "http://vital-iot.org/contexts/service.jsonld",

  "uri": "http://www.example.com#parking",
  "type": "msm:Service",
  "name": "Parking Space Management Service",
  "sawSDL:modelReference":
    "http://www.ibb.gov.tr/ont/ParkingSpaceService",
  "msm:hasOperation": {
    "uri": "http://www.example.com#getParkingStatus",
    "type": "msm:Operation",
    "hrest:hasAddress": "http://www.ibb.gov.tr/parking/{id}",
    "sawSDL:modelReference": "http://www.ibb.gov.tr/ont/ParkingSpace",
    "hrest:hasMethod": "hrest:GET"
  },
  "msm:hasOutput": {
    "uri": "www.example.com#getParkingSpaceStatusReply",
    "type": "wsl:Message",
    "sawSDL:modelReference":
      "http://www.ibb.gov.tr/ont/ParkingSpace",
  }
}
```

Figure 12: Example IoT Service for Parking Space Management (JSON-LD)

Figure 12 shows the same example again, this time in JSON-LD format. Note that we could omit the `uri` and `type` entries to achieve a more compact description. However, we did not do this to make it easier to compare the two representations of the example.

```

{
  "@context": {
    "@vocab": "http://vital-iot.com/ontology#",
    "uri": "@id",
    "vital": "http://vital-iot.com/ontology#",
    "rdfs": "http://www.w3.org/2000/01/rdf-schema#",
    "msm": "http://iserve.kmi.open.ac.uk/ns/msm#",
    "hrest": "http://www.wsmo.org/ns/hrests#",
    "hrest:hasAddress": {
      "@type": "hrest:URITemplate"
    },
    "type": "@type",
    "status": {
      "@type": "@id"
    },
    "status": {
      "@type": "@id"
    },
    "hasOutput": {
      "@type": "@id"
    },
    "hasInput": {
      "@type": "@id"
    },
    "operator": {
      "@type": "@id"
    },
    "serviceArea": {
      "@type": "@id"
    },
    "name": "rdfs:label",
    "description": "rdfs:comment"
  }
}

```

Figure 13: JSON-LD Context Specification for Services

Using this combination of commonly used ontologies, IoT systems can specify their generic services towards the VITAL platform as well as to client applications. The VITAL discovery service can use this information to allow clients to discover IoT services at runtime. The VITAL orchestration service can use it to provide semi-automatic service orchestrations or to provide users with lists of equivalent IoT services. Clearly, this IoT service model can be extended, e.g. with the ability to specify other communication protocols. The consortium will monitor any upcoming need for extensions and address it in an upcoming iteration of this document.

5.3 Conclusion

One of the central aspects of the VITAL system is to integrate existing (background) IoT systems. To do so, we require semantic data models to describe such IoT systems, including metadata about the systems, their operational state and the services that they provide. Note that VITAL does not allow a system to directly specify what sensor data or ICOs it provides – nor is it possible to include actual sensor measurements into the description. Instead, every aspect of the system is

described as an IoT service. To offer access to sensor measurements, a system can specify that it provides access to a data service. To allow pull based data access, a system can specify that it offers a stream management service. This greatly simplifies the system description and reduces its changing rate, while providing a very flexible way to describe (and integrate) systems. The actual interfaces of services are not yet fixed. This is the focus of deliverable D3.2.1, which will be available in M15. However, the general design of VITAL is based on Web technologies. Services are (so far) always RESTful HTTP services. To allow continuous (stream) access, technologies such as Server Side Events or WebSockets will be integrated into VITAL.

In the following we describe the last area that must be modelled in VITAL, the VITAL system itself as well as its services.

6 VITAL SYSTEM AND SERVICES

In addition to (meta) data and services of background IoT systems that are integrated in VITAL, we also need to model the VITAL system itself. This is the aim of this chapter. Note that the chapter does not specify the exact interfaces of systems or services (e.g. the PPI). Instead, its main goal is to discuss data items that are required by these interfaces. Interfaces will be specified in D3.2.1 in month 15.

6.1 VITAL System

A VITAL system is modelled as a collection of RESTful services and an associated semantic metadata description. From the viewpoint of its semantic description, a VITAL system is similar to an IoT system, discussed before. A VITAL system is therefore modelled as a subclass of `IotSystem` with an additional property.

6.1.1 VitalSystem

A `VitalSystem` is defined as a subclass of `IotSystem` that supports an additional VITAL-specific property `providesSystem`. The property `providesSystem` is a subproperty of `ssn:hasSubSystem` that points to a set of `lotSystems`, which are registered in this `VitalSystem`. Note that since a `VitalSystem` is a subset of `IotSystem`, `providesSystem` can also point to another VITAL system, making it possible to create hierarchical VITAL systems.

Class: VitalSystem

SubClassOf: `IotSystem`

ObjectProperty: providesSystem

SubPropertyOf: `ssn:hasSubSystem`

Domain: `VitalSystem`

Range: `IotSystem`

In the following we describe a short example of a VITAL system description in JSON-LD format (see

Figure 14). Note that the example is similar to the example for an IoT system given in Figure 8. Due to the similarities the example uses the same generic external context specification for systems (see Figure 9) as the IoT system example.

```

{
  "@context": "http://vital-iot.org/contexts/system.jsonld",
  "name": "Camden VITAL System",
  "description": "This is the deployment of VITAL in Camden",
  "uri": "http://www.ctu.com/vital/",
  "status": {
    "current": "vital:Running",
    "uri": "http://www.ctu.com/vital/status"
  },
  "operator": "http://www.ctu.com/",
  "serviceArea": "http://dbpedia.org/page/Camden_Town",
  "providesSystem": [
    {
      "uri": "http://www.reply.com/camden"
    },
    {
      "uri": "http://www.ctu.com/xively"
    }
  ],
  "providesService": [
    {
      "type": "StreamManager",
      "msm:hasOperation": [
        {
          "type": "RegisterStream",
          "hrest:hasAddress": "http://www.ctu.com/vital/streamMng/",
          "hrest:hasMethod": "hrest:POST"
        },
        {
          "type": "DeregisterStream",
          "hrest:hasAddress":
            "http://www.ctu.com/vital/stream/{id}/deregister/",
          "hrest:hasMethod": "hrest:GET"
        }
      ]
    },
    {
      "type": "IcoAccess",
      "msm:hasOperation": {
        "hrest:hasAddress": "http://www.ctu.com/vital/ico/{id}",
        "hrest:hasMethod": "hrest:GET"
      }
    }
  ]
}

```

Figure 14: Example VITAL Deployment Description

In the example VITAL system, the description first associates a name (“Camden VITAL System”) and a description (stating that this is a VITAL deployment for Camden Town) with this system. The example system is operated by CTU (as specified by operator). The status is set to Running. Similar to the example for an IoT system, the serviceArea is identified by providing its DBPedia URI, allowing a client to query the DBPedia for more information about Camden. The example system description further specifies that this VITAL instance contains two background IoT systems. The first is an instance of the Reply H1 system (<http://www.reply.com/camden>) the second is the Xively online system (<http://www.ctu.com/xively>). Note that the provided URI is not intended as the actual URI to communicate directly with the background IoT system, e.g.

`https://api.xively.com/` for Xively. Instead the URI provides access to the description of the IoT system (of type `IoTSystem`), further specifying how to access the system (more specifically its adapter classes implementing the VITAL PPI). Finally, the VITAL system description specifies two VITAL services, similar to the ones presented for our example IoT system before (`StreamManager` and `IcoAccess`). Please note that the exact interfaces of VITAL services are not yet defined. Therefore, the examples given here may change in future versions of this deliverable.

In the following we describe the currently planned VITAL services. We also provide models for data required by these services if possible. Since VITAL services are currently under development, the exact data that they need is not always clear. Therefore, this part of the deliverable will be clarified and expanded in future versions. For now the information given should be understood as an outlook and basis for discussions in the consortium.

6.2 Data Access

Data access is one of the basic services of any IoT system, including VITAL. VITAL will provide two different service types with different data access patterns. The first one is a pull-based data access service. The second one is a push-based linked data stream service. For pull-based access, a client specifies the system entity, e.g. the ICO or sensor reading type that it wants to retrieve using its URI.

As an example, to retrieve a temperature reading from sensor `http://www.example.com/vital/sensor/123` in Figure 5, a client would contact the data access service (`IcoAccess`) in

Figure 14 with the URI `http://www.example.com/vital/sensor/123/temperature` and would retrieve the last temperature measurement of this sensor, described in JSON-LD format similar to Figure 6.

Similarly, if a client wants to get continuous updates about the temperature measurements of sensor `http://www.example.com/vital/sensor/123`, then it would contact the `StreamManager` service operation `RegisterStream` at `http://www.ctu.com/vital/streamMng/`, provide the URI of the required data item (again `http://www.example.com/vital/sensor/123/temperature`) and would be registered at the service for a Server Side Event stream that sends new sensor measurements (again using the format shown in Figure 6) when they are created.

As discussed before, the exact interfaces that these services offer are not yet finalised. However, the above discussion shows that the data access service(s) so far do not require additional data models.

6.3 Security and Access Control

6.3.1 Access Control

There are many ontologies available for access control, each utilizing different approaches that provide unique benefits. We evaluated a number of these based on maturity and features offered.

WebAccessControl [WAC] is an ontology created by the Read Write Web Community Group that allows giving access to serialized RDF documents to users and groups, identifying each as HTTP URLs. The main issue with WAC is that it was designed to grant access to the entire document, which makes it unsuited to platforms where more fine-grained permissions are required.

PPO [PPO] and TAC [TAC] are both ontologies designed to extend WAC and allow defining access on a triple-by-triple basis. Which these are more powerful than WAC and allow setting access for each triple, they do not define what exactly access is beyond a Boolean-type value. Both of these ontologies could potentially be extended to better serve VITAL's usage scenario.

S4AC¹⁰ [S4AC] was developed in response to WAC's shortcomings like the two aforementioned ontologies. However, S4AC allows setting fine-grained create, read, update and delete permissions for each data item. This is highly compatible with the permissions model used for traditional REST APIs.

After evaluating all of the above vocabularies we decided to use S4AC for VITAL. The main reason for this decision is that S4AC was the only one to use the CRUD approach for permissions. This kind of fine-grained control is important for a system like VITAL where much of the information is highly confidential.

6.3.2 Users and Authentication in VITAL

Users in VITAL are represented as an extension of the Friend of a friend ontology [FOAF] as well as the fields given in Table 2.

Prefix	Description	Type
vital:userID	ID used to represent user.	Integer
vital:groupOwners	Collection of URLs for users and groups with permission to add and remove users within that group.	Collection of Integers

Table 2: Fields for Modelling Users and Groups

Figure 15 shows some examples for modelling users and groups in VITAL. The first two examples show definitions for two users. The third example shows defining a group that contains the administrators, in this case user 999, for a set of sensors.

¹⁰ http://ns.inria.fr/s4ac/v2/s4ac_v2.html

```
{
  "@context": "http://vital-iot.org/contexts/user.jsonld",
  "@id": "http://www.example.com/vital/users/34423",
  "vital:userID": 34423,
  "@type": "foaf:Person",
  "foaf:name": "Joe Bloggs",
  "foaf:mbox": "joe@bloggs.com"
}

{
  "@context": "http://vital-iot.org/contexts/user.jsonld",
  "@id": "http://www.example.com/vital/users/999",
  "vital:userID": 999,
  "@type": "foaf:Person",
  "foaf:name": "Dave Barlow",
  "foaf:mbox": "dave@barlow.com"
}

{
  "@context": "http://vital-iot.org/contexts/user.jsonld",
  "@id": "http://www.example.com/vital/users/23",
  "vital:userID": 23,
  "@type": "foaf:Group",
  "foaf:name": "Sensor 100-200 Administrators",
  "foaf:mbox": "vital-support@example.com",
  "foaf:member": [
    "http://www.example.com/vital/users/999" ]
}
```

Figure 15: Examples of Users and Groups in VITAL

Users and groups are identified by their type, which is `foaf:Person` for users and `foaf:Group` for groups. Groups may also contain `vital:groupOwners` and `foaf:member` which are both collections of URIS pointing to users or groups. Any field that expects a URI that points to a user in VITAL can also use a URI for a group and vice-versa.

A new VITAL system will always contain a group with `foaf:name` "Administrators" and `foaf:member` with a URI pointing to the initial user who setup the system. The members of this group have full access to the entire system. By creating new groups associated with access control policies defined in section 6.3.3 a VITAL administrator can easily create fine-grained permissions for members.

6.3.3 Integrated Security Model

VITAL uses a system similar to an Access Control List [ACL], using resources URIs as keys. When a resource is requested via a REST endpoint, the URI of that resource along with the user's ID and operation type (using CRUD) is sent to the access control server after the user's identity has been verified. The server then checks if the user's ID is present in the list of authorized users and groups and if the operation type is authorized. Finally, the server responds with a success or fail.

If the requested URI has no permissions associated with it the server will assume only members of the Administrators group can access it. If the resource contains the `foaf:member` property but not the `s4ac:hasAccessPrivilege` property the members of the User Set will be granted Read-Only permissions. In the reverse case, defining permissions without adding members to the members field has no effect.

```
{
  "@id": "http://www.example.com/vital/access-control/patterns/34",
  "pattern": "http://www.example.com/vital/sensor/123/obsvn/*",
  "vital:hasAccessPolicy": [
    {
      "@type": "s4ac:AccessPolicy",
      "s4ac:hasAccessPrivilege": [ "s4ac:Read" ],
      "s4ac:hasAccessConditionSet": {
        "s4ac:hasAccessCondition": [
          {
            "foaf:member": [
              "http://www.example.com/vital/users/34423",
              "http://www.example.com/vital/users/23",
              "http://www.example.com/vital/users/500"
            ]
          }
        ]
      }
    }
  ],
  {
    "@type": "s4ac:AccessPolicy",
    "s4ac:hasAccessPrivilege": [
      "s4ac:Create", "s4ac:Read", "s4ac:Update", "s4ac>Delete"
    ],
    "s4ac:hasAccessConditionSet": {
      "s4ac:hasAccessCondition": [
        {
          "foaf:member": [
            "http://www.example.com/vital/users/457"
          ]
        }
      ]
    }
  }
]
}
```

Figure 16: Access Control Policy for Single Sensor's Observations

Figure 16 describes the access control policy for all observations of Sensor 123. The pattern property describes the pattern that should match this access control policy. While the access control system has not yet been implemented this example is based off the assumption it will utilize a pattern-matching solution that always picks the most specific match for the provided URL.

In this example, the users or groups with the IDs 34423, 23 and 500 have read-only access to sensor 123's observations, while 457 has full access.

6.4 Trust

Trust is an important topic with many different approaches. Some approaches cover trust in terms of trust in entities like sources, contents or agents. Others handle with trust in services and trust based on reliability.

[Sac13] provides the Trust Assertion Ontology (TAO), which describes subjective trust values of asserted users. TAO introduces three classes: `TrustAssertion`, `TrustScale` and `TrustValue`. The social factor based subjective trust values for a source, content or requester are described in `TrustAssertion`. `TrustScale` provides the min and max range for these values. The values themselves within the ranges are specified in `TrustValue`.

[HuFo06] introduces formal semantics and transitivity in terms of an ontology of trust. The concept is based on the idea that trust is built between two entities while each has a finite amount of direct trust relationships. For trust propagation in social networks trust has to be transitive. An example would be if Alice trusts Bob and Bob trusts Eve, Alice would trust Eve as well. Basically [HuFo06] defines two kinds of trust: trust in belief and trust in performance. Typical classes of this ontology are `Trust`, `Trust_b`, `Trust_p` and `Trusted_certainty_degree`.

In VITAL, trust is not restricted to users but does also include trust in (IoT) services and sensors. A number of approaches focus on this kind of trust, e.g. [SNH+10], [HCD06], and [Gal06].

According to [SNH+10], trust is based on expectancy, belief and willingness to be vulnerable for this belief. It is based on three general trust concepts: trust bootstrapping, trust composition and trust propagation. The last two are specific to services. The main concept of the trust service ontology is to provide not just a service trust but also trust as a service. [SNH+10] focuses on reputation-based trust in a service-oriented Web and is based on the notion of community. Trust, trust service and operation are the basic concepts.

In [HCD06] a trusting agent can trust in an agent, a product or a service. Trust in a service is understood as the belief that the service will provide the quality that is expected by the trusting agent. A service trust is measured by some criteria, which should be agreed by the service requester as well as the service provider (both part of the service trust).

In addition to trust in semantic web services, [Gal06] also covers the aspect of *reliability*. Trust is based on ability and reliability. Clients and Web Services are able to describe their trust policies in two ways. First, they can display their own guarantees like security certifications or execution parameters. Second, they can announce their trust preferences in terms of communication partners. [Gal06] handles the problem of trust-based selections as a classification problem. The class

of selected services will contain all services that fit to the trust criteria of the client as well as to the criteria of the service. Used classes are `Goal`, `User` and `WebService`. `Goal` is a class because according to [Gal06] trust perceptions depend on the goal of the requester.

It is currently very difficult to select the best ontology for trust in VITAL. Clearly, the ability to model trust between services is highly relevant to VITAL, making approaches that are able to model this kind of trust (like [SNH+10], [HCD06], and [Gal06]) very interesting. On the other hand, VITAL aims at using lightweight ontologies if possible, for which [Sac13] and [HuFo06] would be better suited. Since no final decision can be taken so far, this is delayed until the second version of this deliverable. As a starting point, we will use [Gal06], since this ontology contains a number of very interesting concepts.

6.5 Discovery and Filtering

The VITAL discovery and filtering services are under active development. Therefore it is not clear which additional data they will need. So far the services concentrate on ICOs, especially sensor devices and need information about a sensor's location (if available), its movement pattern (if known) and its network connectivity (i.e. whether it is connected intermittently or continuously). Therefore, we extend a sensor description (`ssn:Sensor`) with the following properties:

- `hasMovementPattern`, a mandatory property that links to an instance of `MovementPattern`,
- `hasNetworkConnection`, an optional property that links to an instance of `NetworkConnection`.
- `hasLocalizer`, an optional property that links to an IoT service specification that provides access to the current location of the sensor.

Note that the location of a sensor is already modelled in VITAL using the `hasLastKnownLocation` property.

ObjectProperty: `hasMovementPattern`

Domain: `ssn:Sensor`

Range: `MovementPattern`

ObjectProperty: `hasNetworkConnection`

SubPropertyOf: `ssn:hasSubSystem`

Domain: `ssn:Sensor`

Range: `NetworkConnection`

ObjectProperty: hasLocalizer

SubPropertyOf: ssn:hasSubSystem

Domain: ssn:Sensor

Range: msm:Service

6.5.1 MovementPattern

A movement pattern specifies how a device is moving or is expected to move. VITAL currently defines three basic movement patterns: *Stationary* (i.e. no movement at all), *Mobile* (i.e. no additional information is known), and *Predicted* (i.e. there is data available to predict a sensor's future mobility). Note that it is possible to include multiple values for a movement pattern, stating e.g. that a sensor movement pattern is mobile and predicted.

Stationary

If a sensor is stationary, its `hasMovementPattern` property points to an instance of *Stationary*. In this case the sensor may also have a `hasLastKnownLocation` property, pointing to its location. If no such property is present, the sensor location is (currently) unknown. Note that this might change, so a client cannot assume that the sensor location will not be known in the future.

Class: Stationary

SubClassOf: MovementPattern

Mobile

A *Mobile* pattern implies that a sensor may change its location dynamically. Note that this does not necessarily imply that the location is known. If the description of a mobile sensor includes a `hasLastKnownLocation` property, then the provided location *may* be out-dated (which is why we call the property 'last known location'). A system *should* update the last known location field in the sensor description but it *can* do so when it chooses. This reduces the load on a system, which otherwise would need to update a sensor description every time a new location reading becomes available. A sensor description *should not* include a last known location if the system providing it does not update it. A mobile sensor *should* include a `hasLocalizer` property in its description, pointing to a localisation service.

Class: Mobile

SubClassOf: MovementPattern

Predicted

If the movement pattern of a sensor is set to `Predicted`, then the system has additional information about the sensors movement pattern to predict future movements and locations.

Class: Predicted

SubClassOf: MovementPattern

It is currently unclear which exact movement pattern data is needed as this depends on the algorithms that will use them. This is therefore left for future work. An example for a possible prediction pattern is given in Figure 17. The example shows a sensor with a predicted movement pattern that is based on linear interpolation using a predicted speed and movement direction.

```
{
  "@context": "http://vital-iot.org/contexts/sensor.jsonld",
  "name": "TemperatureSensor No.123",
  "type": "VitalSensor",
  "description": "This is an example sensor",
  "uri": "http://www.example.com/vital/sensor/123",
  "hasMovementPattern": {
    "type": "Predicted",
    "hasPredictedSpeed": {
      "value": "3.1",
      "qudt:unit": "qudt:KilometerPerHour"
    },
    "hasPredictedDirection": {
      "type": "NormalVector",
      "geo:lat": "53.2719",
      "geo:long": "-9.0489"
    },
  },
}
,
"ssn:observes": {
  "type": "http://lsm.der.iie/OpenIoT/Temperature",
  "uri": "http://www.example.com/vital/sensor/123/temperature"
}
}
```

Figure 17: Example Sensor Description with Predicted Movement Pattern

6.5.2 NetworkConnection

The `NetworkConnection` specifies how a sensor is connected at the moment, e.g. whether the connection is stable or not. It does not necessarily provide a fully detailed description of the network quality of service that is available. However, it always specifies the expected connection stability. To realise this, so far `NetworkConnection` supports two properties:

- **hasStability**, a mandatory property linking to a `ConnectionStability` instance,
- **hasNetworkSupport**, an optional property that is a sub property of `net:networkSupport` taken from the Network part of the DC ontology [FoLe09], linking to an instance of `net:NetworkSupport` from the same ontology.

ObjectProperty: hasStabilityDomain: `NetworkConnection`Range: `ConnectionStability`**ObjectProperty: hasNetworkSupport**SubPropertyOf `net:networkSupport`Domain: `ssn:Sensor`Range: `net:NetworkSupport`

An example for these properties is given in Figure 18. This example also contains an example for `ConnectionStability`. The described sensor has a stable, continuous network connection and is connected to a single, wired network. More information about this network could be added, e.g. the bandwidth, etc.

```
{
  "@context": "http://vital-iot.org/contexts/sensor.jsonld",
  "name": "TemperatureSensor No.123",
  "type": "VitalSensor",
  "description": "This is an example sensor",
  "uri": "http://www.example.com/vital/sensor/123",
  "hasMovementPattern": {
    "type": "Stationary",
  },
  "hasNetworkConnection": {
    "hasStability": {
      "type": "Continuous"
    },
    "hasNetworkSupport": {
      "net:connectedNetworks": {
        "type": "net:WiredNetwork"
      }
    }
  }
}
```

Figure 18: Example Sensor Description with Network Connection Data

6.5.3 ConnectionStability

`ConnectionStability` allows specifying whether a sensor is connected to a communication network continuously, intermittently or not at all (which will usually not happen). To do so we define three subclasses `Continuous`, `Intermittent` and `Disconnected`. An example for a stable connection is given in Figure 18 above.

Class: Continuous

SubClassOf: `ConnectionStability`

Class: Intermittent

SubClassOf: `ConnectionStability`

Class: Disconnected

SubClassOf: `ConnectionStability`

6.5.4 Localizer

The property `hasLocalizer` is used to refer to a *localizer service*. Such a service provides (REST-full) access to the current location of a sensor, which is especially important for mobile sensors. The service is modelled as an IoT Service as specified in Section 5.2. Note that the localizer service can be realised in different ways, depending on the sensor at hand. As an example, the sensor could provide the service itself, e.g. by giving access to its local GPS receiver. Alternatively the sensor could be localised using an external tracking system and the localizer service could be provided by the IoT system that manages the sensor. An example for a sensor-based localizer service is given in Figure 19.

ObjectProperty: hasLocalizer

Domain: `ssn:Sensor`

Range: `msm:Service`

6.5.5 Queries

In addition to all data items described before, the discovery and filtering services require a way to specify a requested data query, either for ICOs in the case of the discovery service or for data in the case of the filtering service. Due to the nature of Linked Data we do not require a new data model for this. Instead VITAL reuses existing query languages that are well known in the Linked Data community. For onetime queries, VITAL reuses SPARQL, which allows specifying complex sub graph patterns. A SPARQL endpoint will be made available as a RESTful service. Similarly, for continuous (streaming) queries, VITAL reuses the CQELS query language [cq], which extends SPARQL with support for data streams and time windows.

```

{
  "@context": "http://vital-iot.org/contexts/sensor.jsonld",
  "name": "TemperatureSensor No.123",
  "type": "VitalSensor",
  "description": "This is an example sensor",
  "uri": "http://www.example.com/vital/sensor/123",
  "hasMovementPattern": {
    "type": "Mobile"
  },
  "hasLocalizer": {
    "type": "GpsService",
    "msm:hasOperation": {
      "type": "GetLocation",
      "hrest:hasMethod": "hrest:GET",
      "hrest:hasAddress":
        "http://www.example.com/vital/sensor/123/location/"
    },
  },
  "ssn:observes": {
    "type": "http://lsm.der.iie/OpenIoT/Temperature",
    "uri": "http://www.example.com/vital/sensor/123/temperature"
  }
}

```

Figure 19: Example Sensor Description With Localizer

6.6 Complex Event Processing

Events can be understood in different ways and thus are modelled on different layers using different ontologies. In general, there are two basic groups of events: physical events and social events. Physical events occur i.e. in terms of sensors. A physical-event description includes the description of the sensed event, its structure, its trigger and the relations it has with other activities. This kind of event is e.g. modelled in the Semantic Sensor Network Ontology (SSN)¹¹, in OAE¹² or in DOLCE-Lite¹³. In social events it is more important to include what is happening where and by whom. The description generally includes an agent, a time and a place. More details can be found in the Event Ontology¹⁴, the DCMI Event¹⁵, the LODÉ¹⁶ Ontology or in the Schema¹⁷ Ontology.

Real-world events and complex event processing (CEP) is an important current research area in IoT and smart city systems. VITAL provides a CEP service as one

¹¹ <http://www.w3.org/2005/Incubator/ssn/ssnx/ssn>.

¹² <http://www.oae-ontology.org/index.php>

¹³ <http://www.loa.istc.cnr.it/ontologies/DOLCE-Lite.owl>

¹⁴ http://motools.sourceforge.net/event/event.html#term_Event.

¹⁵ <http://dublincore.org/documents/dcmi-terms/#dcmitype-Event>.

¹⁶ <http://linkedevents.org/ontology/>.

¹⁷ <http://schema.org/Event>.

of its add-value services on top of its unified, platform-agnostic view on federated IoT systems. The resulting data models to specify event requests and responses as well as the functionality of a CEP service itself can be rather detailed and complex. After a study of previous work and related approaches, the VITAL consortium agreed on partnering with CityPulse, a current FP7 project in which NUI Galway is participating as a partner. (Complex) event detection and processing is one of the core topics of CityPulse and thus receives a large amount of attention and work in it. VITAL will reuse the semantic models for events, event patterns, requests, responses as well as CEP services that are developed in CityPulse at the moment and will be described in deliverable D3.2 (“Data Federation and Aggregation in Large-Scale Urban Data streams”) of CityPulse. Since VITAL will work on CEP in parallel to this, we will use the current CityPulse model (see, e.g. [GCB14]). The resulting strategic partnership between both project will benefit CityPulse by widening the user base for their event ontology as well as VITAL by ensuring that the project has access to a mature event ontology that extends far beyond the time and resource capabilities of VITAL itself.

6.7 Monitoring

The VITAL monitoring service allows the retrieval of detailed information about a system deployment, a service or an ICO. To do so, the monitoring service provides access to a collection of state updates as measured by the services or ICOs, as well as more detailed descriptions of sensor hardware and software.

6.7.1 State Updates

State updates of the monitoring service (e.g. related to performance) are modelled as virtual sensor measurements. This makes it possible to reuse much of the data models defined before in Chapter 3.4. To retrieve these virtual sensor measurements, clients can contact RESTful interfaces, similar to the ones used for normal measurements. Using standard REST and SSE can support both onetime and streaming updates. The monitoring service can provide information about components like sensors, services and systems. The first version of VITAL supports a core set of monitoring updates that can be extended as required. For each monitoring update we specify a new type of observation as follows:

- **SysUptime**, for a measurement of the uptime of a component,
- **SysLoad**, for a measurement of the load (CPU) of the component,
- **UsedMem**, for a measurement of the used memory of a component,
- **AvailableMem**, for a measurement of the available memory of a component,
- **ServedRequests**, for a measurement of the total number of requests that a component has served between the last re-start and the time of the measurement,
- **ActiveRequests**, for a measurement of the number of requests that were served at the time of the measurement,
- **MaxRequests**, for a measurement of the maximum number of requests that a component can serve simultaneously,
- **Errors**, for a measurement of the total number of errors that have occurred between the last re-start and the time of the measurement.

In addition, we reuse the class `OperationalState` (specified in Section 5.1.2) to specify a measurement about the current operational state of a system or service. `OperationalState` will be extended with new states depending on application requirements (possible new states are `LostConnection`, `InError`, `LowBattery`).

Figure 20 shows an example of a virtual measurement send by the monitoring service. The measurement specifies the uptime of sensor with the identifier `http://www.example.com/vital/sensor/123` as 2023546 milliseconds (approximately 34 minutes). Other measurements can be constructed similarly.

```
{
  "@context": "http://vital-iot.org/contexts/measurement.jsonld",
  "uri": "http://www.example.com/vital/sensor/123/obsvn/42",
  "type": "ssn:Observation",
  "ssn:observationProperty": {
    "type": "Uptime"
  },
  "ssn:observationResultTime": {
    "inXSDDateTime": "2014-08-23T14:03:11+01:00"
  },
  "observationSubject": "http://www.example.com/vital/sensor/123",
  "ssn:observationResult": {
    "type": "ssn:SensorOutput",
    "ssn:hasValue": {
      "type": "ssn:ObservationValue",
      "value": "2023546",
      "qudt:unit": "qudt:MilliSecond"
    }
  }
}
```

Figure 20: Example Virtual Measurement for System Uptime

6.7.2 Activity Logging

In addition to updates about the state of a system component, the monitoring service may also provide updates about activities occurring in the system. These updates could be stored in an activity log, e.g. to log a user logging in or a system starting up. Such updates can be modelled similarly to state updates as virtual measurements by specifying new observation types.

6.7.3 Sensor Hardware and Software

Since VITAL reuses the DC ontologies, it already includes a detailed data model for hardware and software components like a battery, a keyboard, an operating system, etc. These models can be used to provide a detailed description of both the components and the status of a sensor. To do so, a sensor description can include the DC ontology properties `hard:deviceHardware` and `soft:deviceSoftware`, pointing to an instance of `hard:DeviceHardware` and `soft:DeviceSoftware` respectively. These may link to additional information about the sensor's hardware and software parts, including their current state (e.g. using `hard:status` to link to

an instance of `hard:HardwareStatus`). An example for this is shown in Figure 21. The example description specifies a sensor that is currently active, has a build in memory size of 128 kByte and a CPU with a maximum speed of 10 MHerz.

```
{
  "@context": "http://vital-iot.org/contexts/sensor.jsonld",
  "name": "TemperatureSensor No.123",
  "type": "VitalSensor",
  "description": "This is an example sensor",
  "uri": "http://www.example.com/vital/sensor/123",
  "deviceHardware": {
    "hard:status": "hard:HardwareStatus_ON",
    "hard:builtInMemory": {
      "size": 131072
    },
    "hard:cpu": {
      "type": "hard:CPU",
      "maxCpuFrequency": 10
    }
  }
}
```

Figure 21: Example Sensor Description with Hardware Components

6.8 Orchestration and Development

No specific data items and models for the orchestration service and the VITAL development environment have been identified so far. The consortium will continue to monitor these areas and extend this deliverable as necessary.

6.9 Conclusion

VITAL handles VITAL system deployments as special cases of IoT system deployments. This allows creation of hierarchical systems that contain several VITAL systems. As an example, a smart city like Istanbul might operate a local installation of VITAL that integrates all IoT systems in Istanbul. In addition, this VITAL system might be integrated into a larger, transnational VITAL system that may e.g. also include London and other cities. Besides easier administration, this approach enables e.g. the ability to fine-tune access rights and security, possibly hiding specific services and ICOs when integrating the system into the transnational one.

As discussed before, VITAL services are under active development, making it difficult to define their data items precisely. In this chapter we provided a first version of data items that have been identified so far as well as the current state of our analysis for data models in different areas, e.g. trust, etc. Based on this the project will work on defining in detail the interfaces between all system components, leading to deliverable D3.2.1.

7 CONCLUSION AND NEXT STEPS

This deliverable provides the basis for the semantic (meta-) data models used in the VITAL system. We build upon Linked Data principles and technologies to provide interoperable and platform agnostic data models that are based on existing ontologies. This allows VITAL applications to integrate other data sources in the Web, resulting in a large and varied set of usable data items. Although we analysed a large number of ontologies during the design of the VITAL data models, the work is not finished. Firstly, for some system components like VITAL services, the actual data needed are not yet fully clear. While these services are developed, new data requirements will arise and have to be addressed. In addition, for some areas, e.g. unit of measurements or trust, it is not entirely clear if the chosen ontologies are always the best choices. With more experience in developing real VITAL applications, we will develop a better understanding of the advantages and disadvantages of these ontologies and may have to adapt some decisions. This is to be expected in the context of a large and complex project like VITAL. Therefore, we will continue to actively monitor the usage of the VITAL data models in the project (and possibly beyond) to identify any need for additional work or ontologies.

As a prominent example for a possible extension in the second iteration of this deliverable, we have so far omitted the modelling of actuators and actuator commands. This is an active field of research in the semantic IoT community and could allow to develop much more powerful Linked Data based IoT systems by allowing to model commands with RDF or to execute them with SPARQL. We will look into this area and will provide a more detailed discussion in D3.1.2. The final version of this deliverable (i.e. D3.1.2) will be integrated in the prototype of the VITAL platform, with a view to cover the needs of the VITAL scenarios and use cases, but also the needs of a wider range of smart cities application scenarios. Additional scenarios may require further enhancements to the VITAL data models, based however on the same thinking and approach employed for deliverable D3.1.1.

Also, based on the current deliverable work in work package WP3 of VITAL will continue with specifying the interfaces between all system components. This will lead to deliverable D3.2.1 as well as a first skeleton implementation of the VITAL system with mock-up components using the real interfaces. Using this skeleton system, all VITAL components and services will then be integrated, leading to a first version of the final VITAL platform.

8 REFERENCES

- [ACL] "Access Control List", http://www.cisco.com/c/en/us/td/docs/ios/12_2/security/configuration/guide/fsecur_c/scfacts.html, last visited 28 August 2014.
- [CBB+12] M. Compton, P. Barnaghi, L. Bermudez, R. Garcia-Castro, O. Corcho, S. Cox, J. Graybeal, M. Hauswirth, C. Henson, A. Herzog, V. Huang, K. Janowicz, W. D. Kelsey, D. Le Phuoc, L. Lefort, M. Leggieri, H. Neuhaus, A. Nikolov, K. Page, A. Passant, A. Sheth, K. Taylor, "The SSN Ontology of the W3C Semantic Sensor Network Incubator", in *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 17, pp. 25-32, Elsevier, 2012.
- [CDN09] A. Caragliu, C. Del Bo, P. Nijkamp, "Smart cities in Europe", in: *Serie Research Memoranda 0048* (VU University Amsterdam, Faculty of Economics, Business Administration and Econometrics), 2009.
- [com] COMPOSE Consortium, "COMPOSE: Collaborative Open Market to Place Objects at your Service Website", <http://www.compose-project.eu/>, last visited 18 August 2014.
- [CoWi01] B. Corona, S. Winter, "Datasets for pedestrian navigation services" in *Angewandte Geographische Informationsverarbeitung, Proceedings of the AGIT Symposium*, 2001
- [CPF+04] H. Chen, F. Perich, T. Finin, and A. Joshi, "Soupa: Standard ontology for ubiquitous and pervasive applications," in *Mobile and Ubiquitous Systems: Networking and Services*, International Conference on, 2004.
- [cq] "Continuous Query Evaluation over Linked Stream (CQELS) – CQELS Language", https://code.google.com/p/cqels/wiki/CQELS_language, last visited 25 August 2014.
- [DSW06] J. Davis, R. Studer, P. Warren, "Semantic Web Technologies: Trends and Research in Ontology based Systems", John Wiley & Sons, 2006.
- [dul] "DOLCE+DnS Ultralite (DUL) – The DOLCE+DnS Ultralite ontology", http://ontologydesignpatterns.org/wiki/Ontology:DOLCE+DnS_Ultralite, last visited 24 August 2014.
- [ELS07] M. Eid, R. Liscano, and A. E. Saddik, "A universal ontology for sensor networks data," in *Computational Intelligence for Measurement Systems and Applications (CIMSAS 2007)*, IEEE International Conference on, 2007, pp. 59–62.
- [FaLa07] J. Farrell, H. Lausen, "Semantic Annotations for WSDL and XML Schema", W3C Recommendation, <http://www.w3.org/TR/sawSDL/>, last visited 20 August 2014, August 2007.
- [FOAF] D. Brickley, L. Miller, Friend of a Friend (FOAF) project (of Semantic Web project), <http://www.foaf-project.org/>, 2000, last visited 28 August 2014.
- [FoLe09] J. M. Cantera Fonseca, R. Lewis, "Delivery Context Ontology", W3C Working Draft 16, <http://www.w3.org/TR/2009/WD-dcontology-20090616/>, last visited 25 August 2014, June 2009.

- [Gal06] S. Galizia, "WSTO: A Classification-Based Ontology for Managing Trust in Semantic Web Services" in Proceedings of the 3rd European Semantic Web Conference (ESWC 2006), pp. 697-711, Budva, Montenegro, June 2006.
- [GaSc14] F. Gandon, G. Schreiber, "RDF 1.1 XML Syntax", W3C Recommendation, <http://www.w3.org/TR/rdf-syntax-grammar/>, last visited 24 August 2014, 25 February 2014.
- [GCB14] F. Gao, E. Curry, S. Bhiri, "Complex Event Service Provision and Composition based on Event Pattern Matchmaking", in Proceedings of DEBS'14, Mumbai, India, May 2014.
- [GoRu06] C. Goodwin and D. J. Russomanno, "An ontology-based sensor network prototype environment," in Information Processing in Sensor Networks, Fifth International Conference on, 2006.
- [GRS10] K. Gomadam, A. Ranabahu, A. Sheth, "SA-REST: Semantic Annotation of Web Resources", W3C Member Submission, <http://www.w3.org/Submission/SA-REST/>, last visited 20 August 2014, 05 April 2010.
- [HaMa12] Mohamed H. Haggag and Doaa R. Mahmoud, "OnTraJaCS: Ontology based Traffic Jam Control System", in International Journal of Computer Applications vol. 60(2), p. 6-16, December 2012.
- [HASB13] I. Herman, B. Adida, M. Sporny, M. Birbeck, "RDFa 1.1 Primer - Second Edition: Rich Structured Data Markup for Web Documents", W3C Working Group Note, <http://www.w3.org/TR/rdfa-primer/>, last visited 24 August 2014, 22 August 2013.
- [HaSe13] S. Harris, A. Seaborne, eds., "SPARQL 1.1 Query Language", W3C Recommendation, <http://www.w3.org/TR/sparql11-query/>, last visited 20 August 2014, 21 March 2013.
- [HCD06] F. K. Hussain, E. Chang, T. S. Dillon, "Trust ontology for service-oriented environment", in Proceedings of the ACS/IEEE International Conference on Computer Systems and Applications (AICCSA), pp. 320-325, Dubai, UAE, 2006.
- [HeBi11] Tom Heath, Christian Bizer, "Linked Data. Envolving the Web into a Global Data Space", Morgan & Claypool, 2011.
- [HKHD09] B. Heitmann, S. Kinsella, C. Hayes, S. Decker, "Implementing Semantic Web Applications: Reference Architecture and Challenges", in Proceedings of the 5th Workshop on Semantic Web Enabled Software Engineering, at the International Semantic Web Conference (ISWC09), http://ceur-ws.org/Vol-524/swese2009_2.pdf, 2009.
- [HKHS14] R. Hodgson, P. J. Keller, J. Hodges, J. Spivak, "QUDT - Quantities, Units, Dimensions and Data Types Ontologies", <http://qudt.org/>, last visited 22 August 2014, March 18, 2014.
- [HKO+10] M. Houda, M. Khemaja, K. Oliveira, M. Abed, "A public transportation ontology to support user travel planning" in Research Challenges in Information Science (RCIS), 2010 Fourth International Conference on, IEEE, May 2010, pp. 127-136, DOI: 10.1109/RCIS.2010.5507372

- [HoPa06] Jerry R. Hobbs, Feng Pan, “Time Ontology in OWL”, W3C Working Draft, <http://www.w3.org/TR/owl-time>, last visited 21 August 2014, 27 September 2006.
- [HuFo06] J. Huang, M. S. Fox, “An ontology of trust: formal semantics and transitivity”, in Proceedings of the 8th international conference on Electronic commerce (ICEC '06): The new e-commerce: innovations for conquering current barriers, obstacles and limitations to conducting successful business on the internet, pp. 259-270, USA, 2006.
- [isrv] “iServe: Where Linked Data Meets Services“, <http://iserve.kmi.open.ac.uk/>, last visited 18 August 2014.
- [jsls] “JSON for Linking Data”, <http://json-ld.org/>, last visited 24 August 2014.
- [KGV08] J. Kopecký, K. Gomadam, T. Vitvar, “hRESTS: an HTML Microformat for Describing RESTful Web Services”, in Proceedings of the 2008 IEEE/WIC/ACM International Conference on Web Intelligence (WI-08), Sydney, Australia, December 2008.
- [KNSP09] R. Krummenacher, B. Norton, E. Simperl, C. Pedrinaci, “SOA4All: Enabling Web-scale Service Economies”, in Proceedings of the 2012 IEEE Sixth International Conference on Semantic Computing (ICSC09), pp. 535-542, Berkeley, CA, USA, 2009
- [KVF09] J. Kopecky, T. Vitvar, D. Fensel, “D3.4.3 MicroWSMO and hRESTS”, Deliverable 3.4.3, EU FP7 Project SOA4All, <http://sweet.kmi.open.ac.uk/pub/microWSMO.pdf>, last visited 20 August 2014, March 2009.
- [LPR05] H. Lausen, A. Polleres, D. Roman, et al., “Web Service Modeling Ontology (WSMO)”, W3C Member Submission, <http://www.w3.org/Submission/WSMO/>, last visited 20 August 2014, 3 June 2005.
- [LWS13] M. Leggieri, C. von der Weth, M. Serrano, “Semantic Representations of Internet-Connected Objects”, EU PF7 Project OpenIoT Deliverable 3.1.2 (D312-130315-v07), March 2013.
- [MBH+04] David Martin, Mark Burstein, Jerry Hobbs, Ora Lassila, Drew McDermott, Sheila McIlraith, Srini Narayanan, Massimo Paolucci, Bijan Parsia, Terry Payne, Evren Sirin, Naveen Srinivasan, Katia Sycara, “OWL-S: Semantic Markup for Web Services”, W3C Member Submission, <http://www.w3.org/Submission/OWL-S/>, last visited 18 August 2014, 22 November 2004.
- [NoMc01] N. F. Noy, D. L. McGuinness, “Ontology Development 101: A Guide to Creating Your First Ontology”, Stanford Knowledge Systems Laboratory Technical Report KSL-01-05 and Stanford Medical Informatics Technical Report SMI-2001-0880, http://protege.stanford.edu/publications/ontology_development/ontology101.pdf, last visited 20 August 2014, March 2001.
- [OM] “Ontology of units of Measure (OM)”, <http://www.wurvoc.org/vocabularies/om-1.6/>, last visited 25 August 2014.
- [OTN] “Ontology of Transportation Networks”, Deliverable A1-D4, Project REVERSE: Reasoning on the Web with Rules and Semantics, originally available at <http://reverse.net/deliverables/m18/a1-d4.pdf>, Unavailable as of 27 August 2014, cached version available at <https://webcache>.

- googleusercontent.com/search?q=cache:gi0h1sBXLHYJ:reverse.net/deliverables/m18/a1-d4.pdf #36, last visited 27 August 2014.
- [PML+10] C. Pedrinaci, D. Liu, M. Maleshkova, D. Lambert, J. Kopecký, J. Domingue, “iServe: a linked services publishing platform”, in: Proceedings Workshop on Ontology Repositories and Editors for the Semantic Web, at the 7th Extended Semantic Web Conference, Heraklion, Greece, June 2010.
- [PoBe08] L. Polo, D. Berrueta, “MUO Measurement Units Ontology”, Working Draft, <http://idi.fundacionctic.org/muo/muo-vocab.html>, last visited 25 August 2014, April 2008.
- [PPO] “Privacy Preference Ontology (PPO)”, <http://vocab.deri.ie/ppo>, last visited 21 August 2014.
- [S4AC] “S4AC Vocabulary Specification 0.2 Namespace Document 6 October 2011”, http://ns.inria.fr/s4ac/v2/s4ac_v2.html, last visited 21 August 2014.
- [Sac13] O. Sacco, “Trust Assertion Ontology”, DERI, NUI Galway, 2013. <http://vocab.deri.ie/tao>, last visited 20 August 2014.
- [SJB+09] C. Stasch, K. Janowicz, A. Bröring, I. Reis, and W. Kuhn, “A Stimulus-Centric Algebraic Approach to Sensors and Observations,” in Proc. 3rd International Conference on GeoSensor Networks (GSN’09), pp. 169–179, 2009.
- [SKD+09] G. Stevenson, S. Knox, S. Dobson, and P. Nixon, “Ontonym: a collection of upper ontologies for developing pervasive systems,” in Proceedings 1st Workshop on Context, Information and Ontologies (CIAO’09), pp. 1–8, 2009.
- [SLK+14] M. Sporny, D. Longley, G. Kellogg, M. Lanthaler, N. Lindstroem, “JSON-LD 1.0 - A JSON-based Serialization for Linked Data”, W3C Recommendation, 2014, <http://www.w3.org/TR/json-ld/>, last visited 21 August 2014.
- [SNH+10] W. Sherchan, S. Nepal, J. Hunklinger, A. Bouguettaya, “A Trust Ontology for Semantic Services”, in Proceedings of the IEEE 7th International Conference on Services Computing (SCC 2010), pp. 313-320, USA, 2010.
- [SoKa13] J. Soldatos, J. Kaldis, “D2.1 – Report on Stakeholder’s and Virtualization Requirements”, VITAL Project Deliverable D2.2, December 2013.
- [SoKa14] J. Soldatos, J. Kaldis, “D2.2 – Reference and Validation Scenarios for IoT Virtualization”, VITAL Project Deliverable D2.2, April 2014.
- [SRA06] M. Strimpakou, I. Roussaki, and M. E. Anagnostou, “A context ontology for pervasive service provision,” in Advanced Information Networking and Applications (AINA 2006), 20th International Conference on, 2006.
- [ssn] The W3C Semantic Sensor Network Incubator Group, “Semantic Sensor Network Ontology”, <http://www.w3.org/2005/Incubator/ssn/ssnx/ssn>, last visited 22 August 2014.

- [TAC] “TripleAccessControl Ontology 0.1 Namespace Document 15 September 2011”, <http://ns.bergnet.org/tac/0.1/triple-access-control.html>, last visited 21 August 2014.
- [UO] “Units of Measurement Ontology”, <http://bioportal.bioontology.org/ontologies/UO>, last visited 25 August 2014.
- [VHM+14] R. Verborgh, A. Harth, M. Maleshkova, S. Stadtmüller, T. Steiner, M. Taheriyan, R. Van de Walle, “Survey of Semantic Description of REST APIs”, book chapter in: REST: Advanced Research Topics and Practical Applications (C. Pautasso, E. Wilde, R. Alarcon), pp. 69-89, Springer New York, 2014.
- [VKVF08] T. Vitvar, J. Kopecky, J. Viskova, D. Fensel, “WSMO-Lite Annotations for Web Services”, in Proceedings of the 5th European Semantic Web Conference, Tenerife, Spain, 2008.
- [w3c] The W3C Semantic Sensor Network Incubator Group, “W3C SSN Incubator Group Review of Sensor and Observation Ontologies”, http://www.w3.org/2005/Incubator/ssn/wiki/Incubator_Report#Review_of_Sensor_and_Observation_ontologies, last visited 20 February 2014.
- [WAC] “WebAccessControl”, <http://www.w3.org/wiki/WebAccessControl>, last visited 21 August 2014.
- [WeBa09] W. Wei and P. Barnaghi, “Semantic annotation and reasoning for sensor data,” in Smart sensing and context (EuroSSC’09), 4th European conference on, 2009, pp. 66–76.