LDAC 2024

Lessons learnt from researching on semantic interoperability

Raúl García-Castro

Ontology Engineering Group
Centro de I+D+i en Inteligencia Artificial (AI.nnovation Space)
Universidad Politécnica de Madrid

rgarcia@fi.upm.es

13/06/2024

Bochum

Semantic
interoperability

Cognitive science

Logic

Knowledge representation

"How to represent facts about the world in a computer-understandable form"

Natural Language Processing

Automated reasoning

Computer vision

Robotics

Rational agents

Machine learning

S. Russel and P. Norvig "Artificial Intelligence, a modern approach". Third Edition. Prentice Hall. 2010.

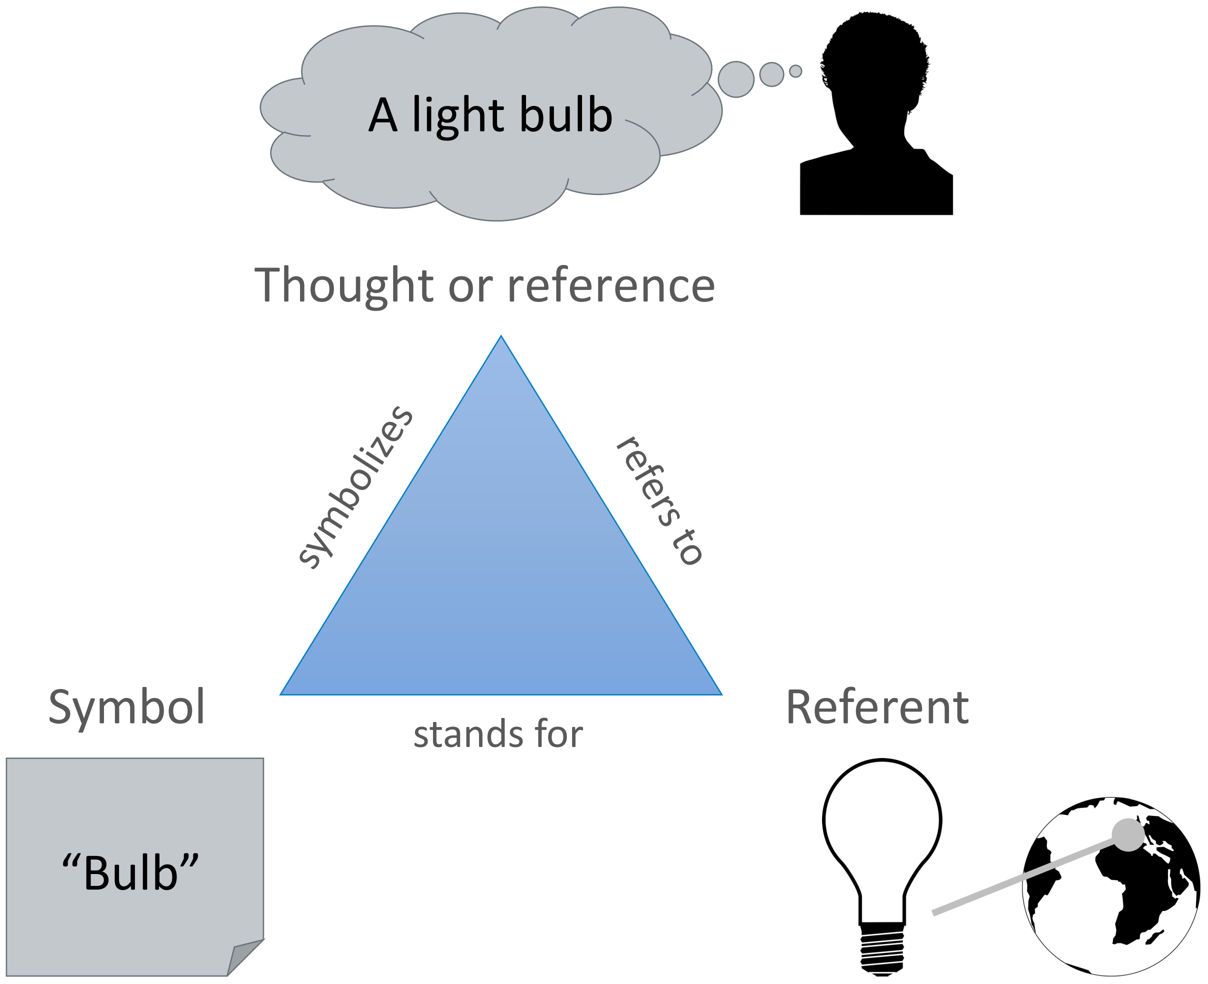

C. K. Ogden and I. A. Richards (1923) The Meaning of Meaning

What do we need to share?

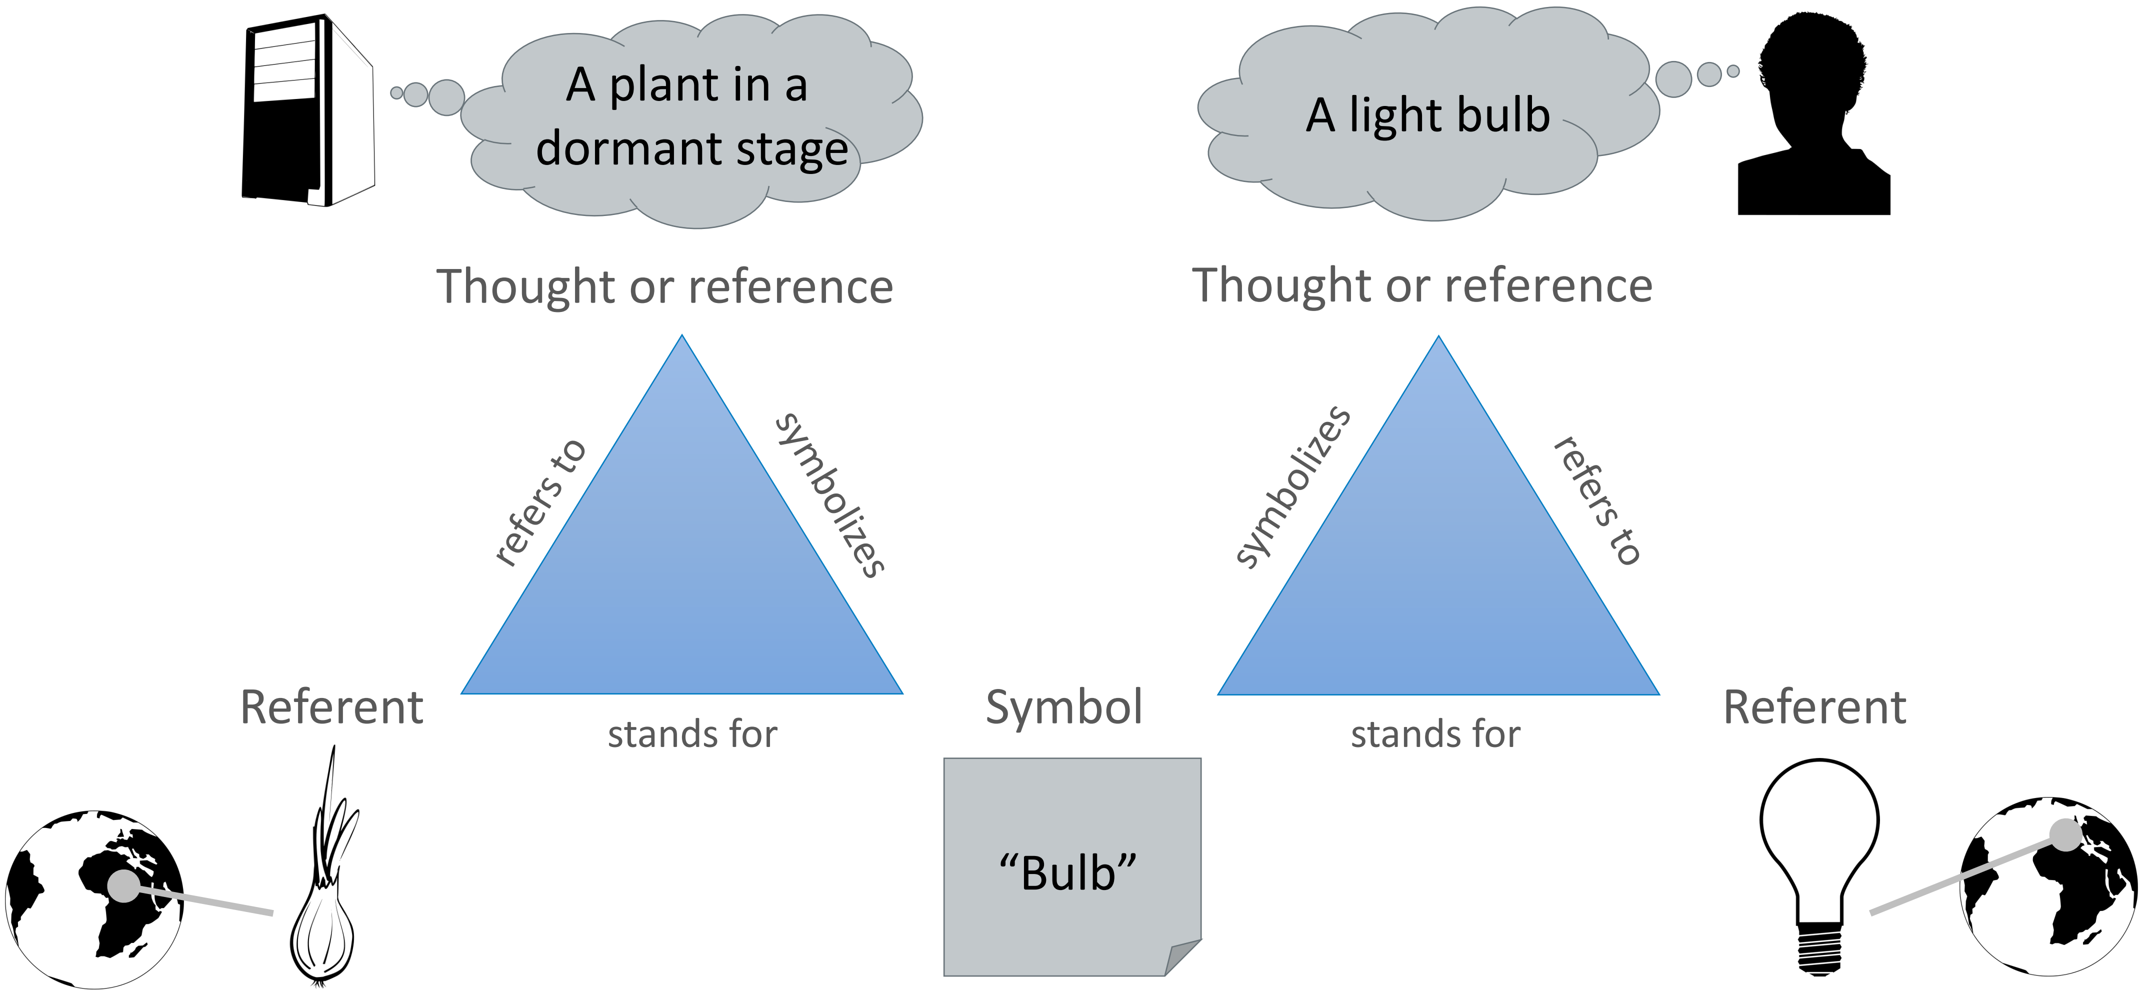

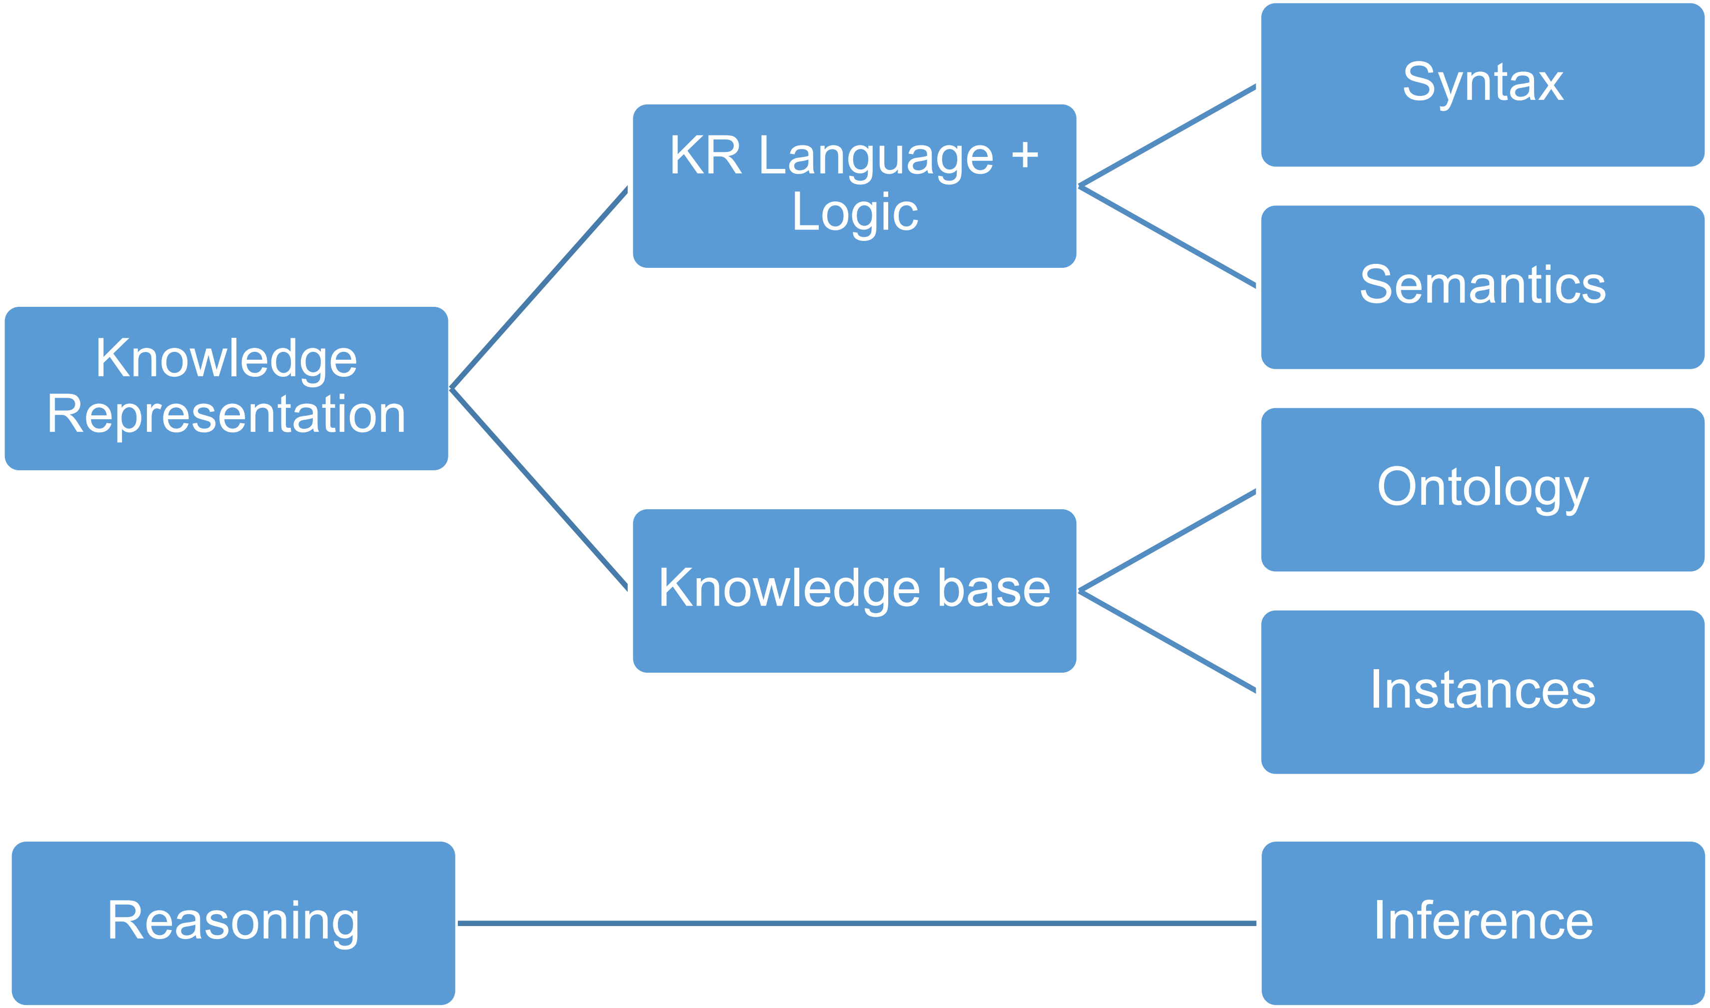

S. Russel and P. Norvig "Artificial Intelligence, a modern approach". Third Edition. Prentice Hall. 2010.

How many semantics?

Am I using correctly the semantics?

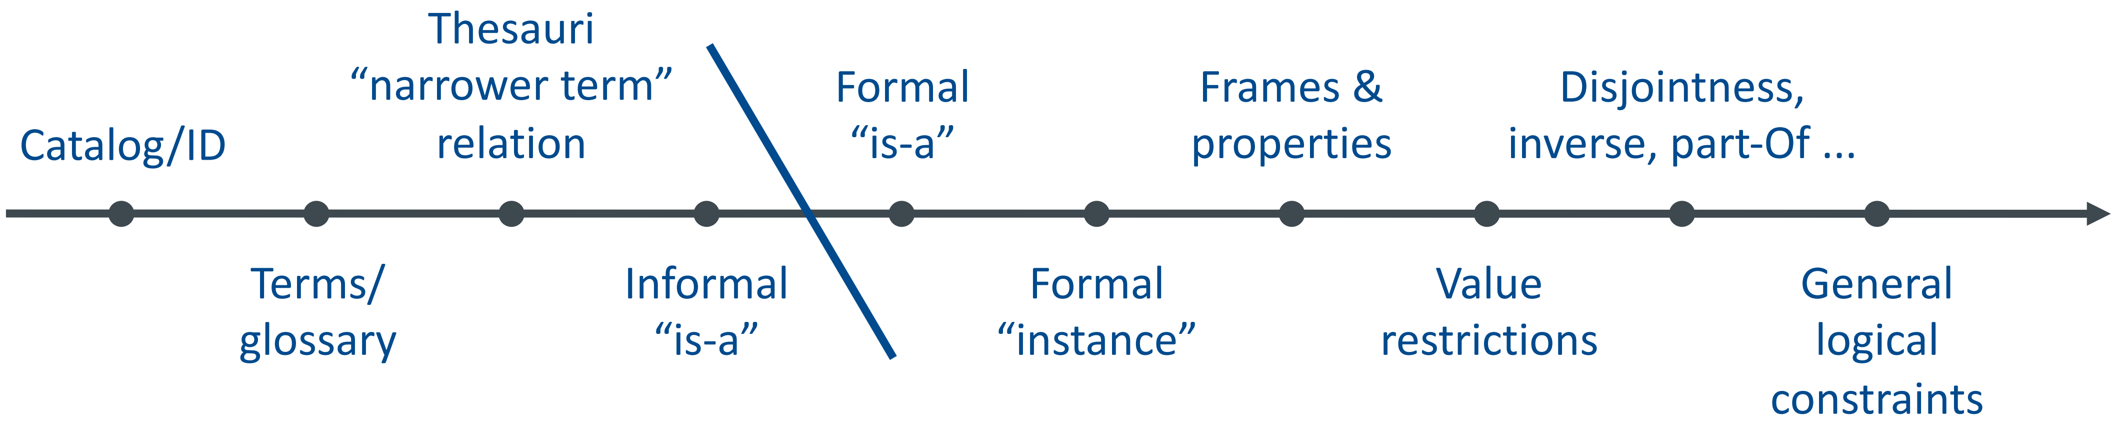

Lassila O, McGuiness D. The Role of Frame-Based Representation on the Semantic Web. Technical Report. Knowledge Systems Laboratory. Stanford University. KSL-01-02. 2001.

Correct concept definitions

Characteristics of a concept:

- **Essential**: indispensable to understand and distinguish a concept
- **Complementary**: colour, material, shape, …

		
Leather	Plastic	Plastic
Red	Grey	Grey
Wheel	Wheel	No Wheel
Store and transport clothes and personal items when traveling	Store and transport clothes and personal items when traveling	Store and transport clothes and personal items when traveling
Resistant material	Resistant material	Resistant material
Handles	Handles	Handles

Every citizen must have a DNI

Global restriction

```
Class: Citizen
Class: DNI
ObjectProperty: hasIdCard
     Characteristics: Functional
     Domain: Citizen
     Range: DNI
```

Local restriction

```
Class: Citizen
     SubClassOf: hasIdCard exactly 1 DNI
Class: DNI
ObjectProperty: hasIdCard
```

Does every citizen have a DNI?

GR: No

LR: Yes

Beware of the assumptions when using an ontology language!

- Open World Assumption
  - The lack of evidence over a fact does not imply that it is false or that it does not hold

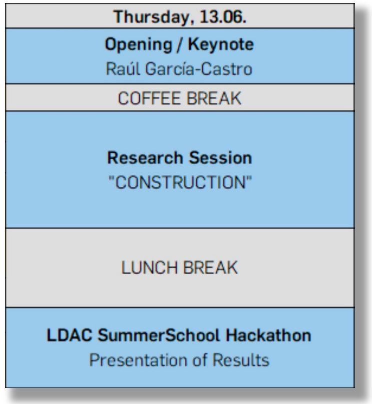

Is there a speaker in the coffee break?

CW: No

OW: I don't know

- Non unique name assumption
  - Different URIs do not necessarily identify different individuals

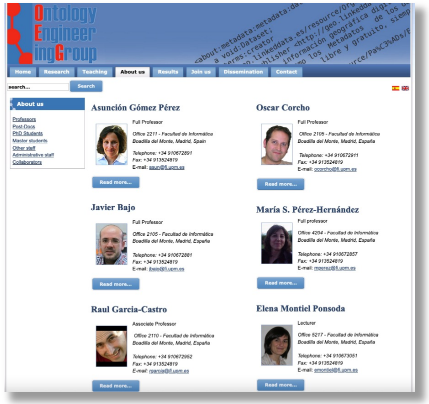

How many members has the OEG?

UNA: 40

NUNA: At least 1

Linked Open Terms – industrial methodology

http://lot.linkeddata.es/

- Towards lightweight and agile processes
- Inspiration from software development practices
- Coupling software and ontology development

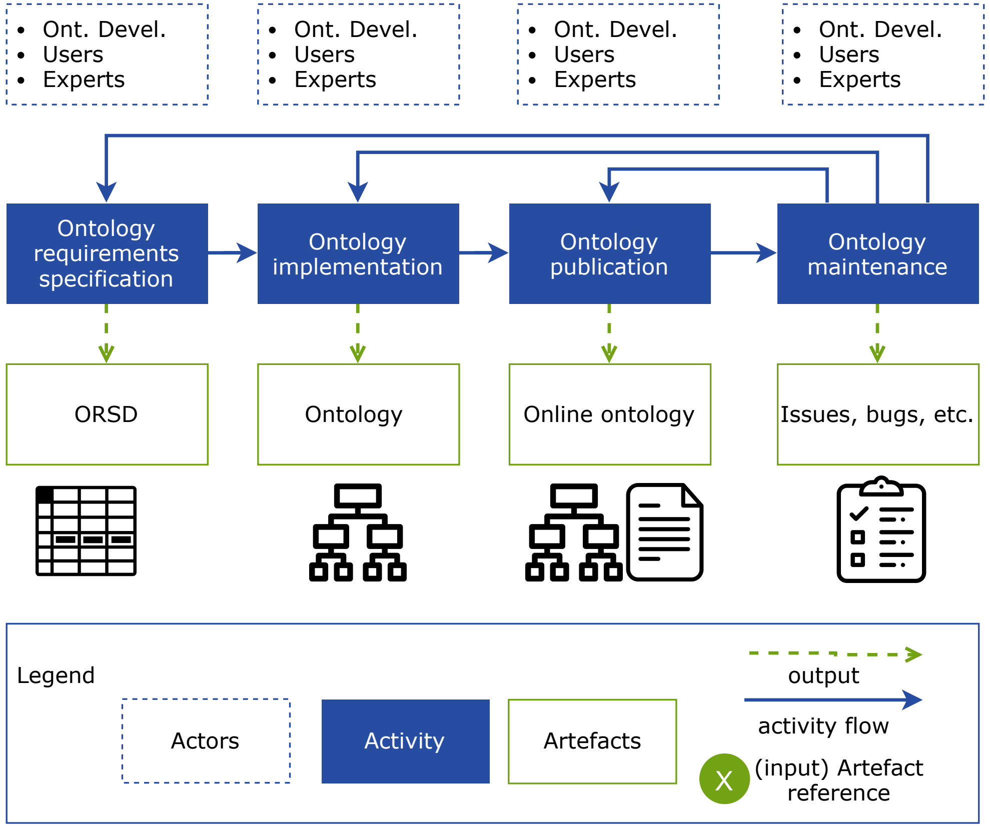

María Poveda-Villalón, Alba Fernández-Izquierdo, Mariano Fernández-López, Raúl García-Castro. LOT: An industrial oriented ontology engineering framework. Engineering Applications of Artificial Intelligence, Volume 111, May 2022, 104755. https://doi.org/10.1016/j.engappai.2022.104755

Many artefacts in the ontology engineering process

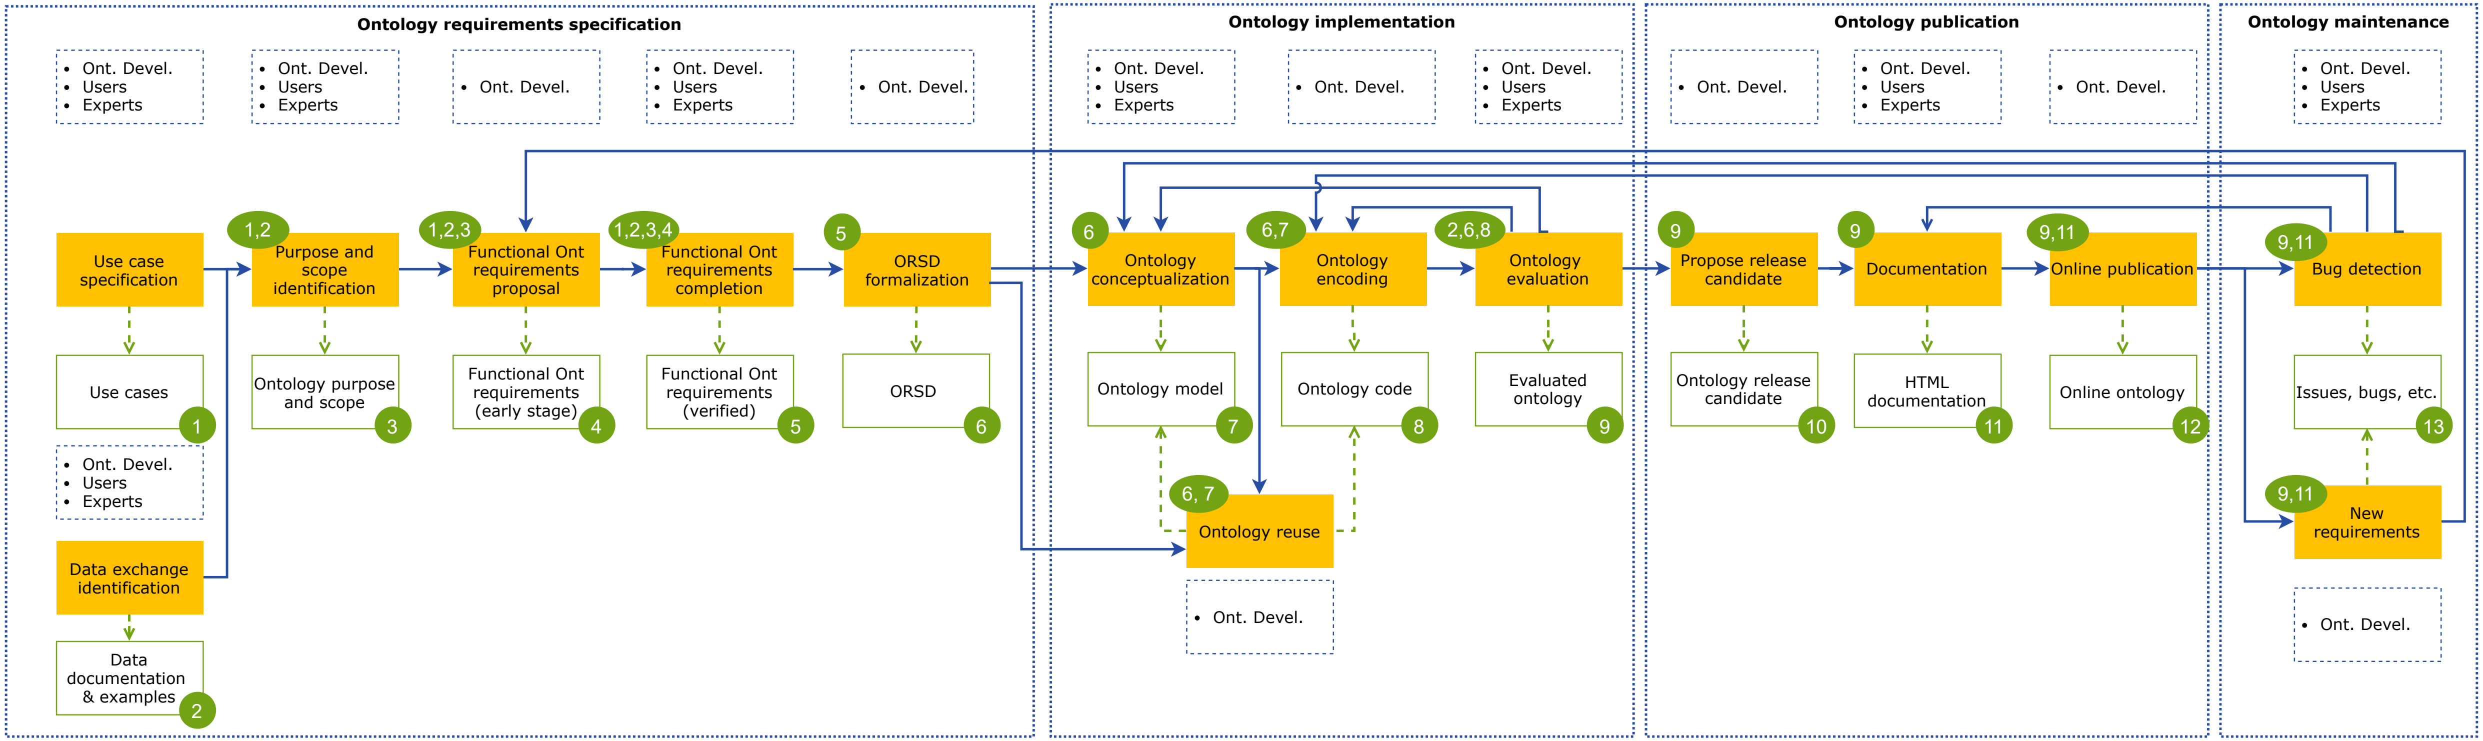

- Ontology engineering artefacts must ensure:
    - Traceability
    - Currentness
    - Consistency
    - etc.

Need to monitor the development process

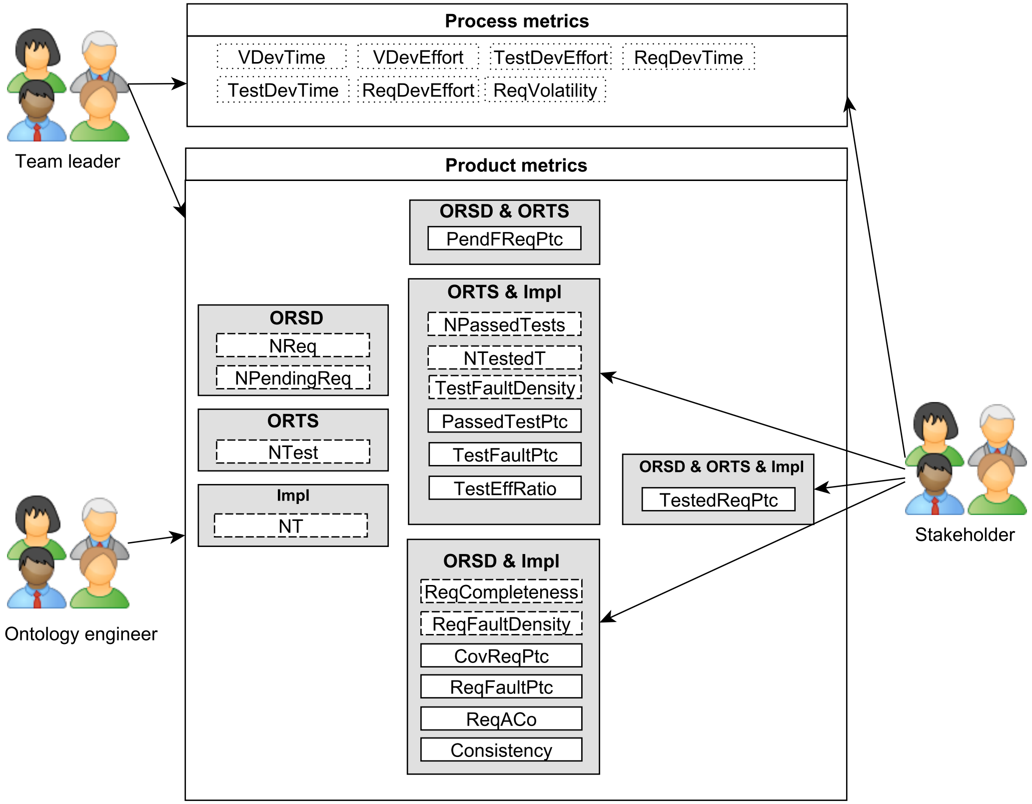

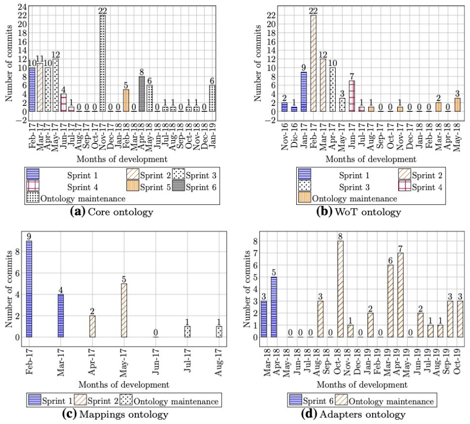

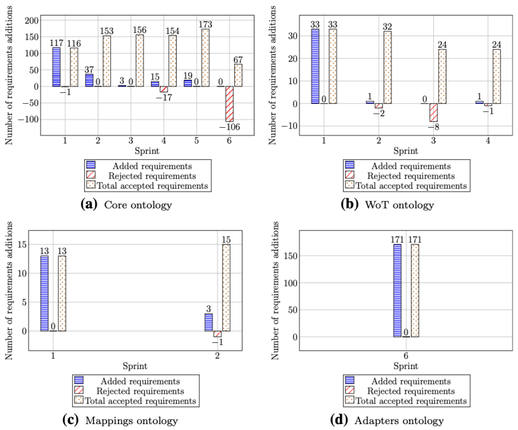

Metrics-driven analysis:

- The number of requirements or their complexity do not influence individually the development time
- However, if the complexity and the number of requirements is high, then the time increases

Computing metrics is time consuming;
it is essential to have software that supports the automatic generation of metrics

Fernández-Izquierdo, A., Poveda-Villalón, M., Gómez-Pérez, A. et al. Towards metrics-driven ontology engineering. Knowl Inf Syst 63, 867–903 (2021). https://doi.org/10.1007/s10115-021-01545-9

Open ontology development artefacts

https://auroral.iot.linkeddata.es/

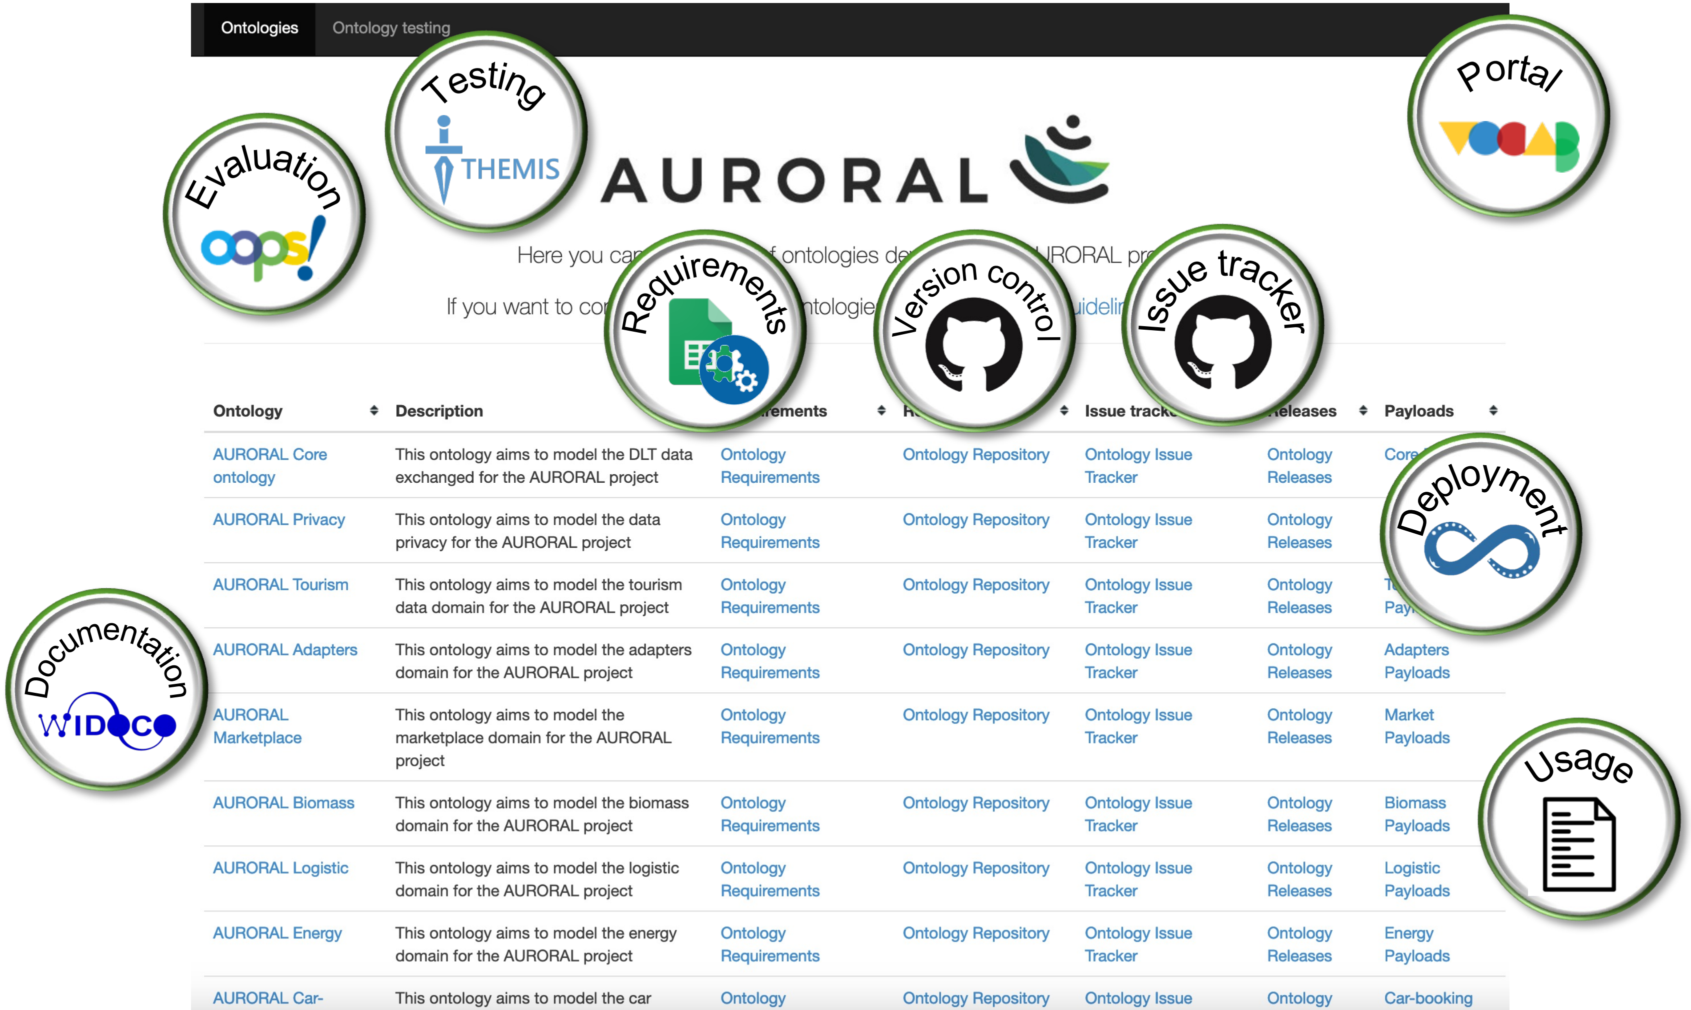

Ontology	Description	rements	R	Issue track	eleases	Payloads
AURORAL Core ontology	This ontology aims to model the DLT data exchanged for the AURORAL project	Ontology Requirements	Ontology Repository	Ontology Issue Tracker	Ontology Releases	Core
AURORAL Privacy	This ontology aims to model the data privacy for the AURORAL project	Ontology Requirements	Ontology Repository	Ontology Issue Tracker	Ontology Releases	
AURORAL Tourism	This ontology aims to model the tourism data domain for the AURORAL project	Ontology Requirements	Ontology Repository	Ontology Issue Tracker	Ontology Releases	T Pay
AURORAL Adapters	This ontology aims to model the adapters domain for the AURORAL project	Ontology Requirements	Ontology Repository	Ontology Issue Tracker	Ontology Releases	Adapters Payloads
AURORAL Marketplace	This ontology aims to model the marketplace domain for the AURORAL project	Ontology Requirements	Ontology Repository	Ontology Issue Tracker	Ontology Releases	Market Payloads
AURORAL Biomass	This ontology aims to model the biomass domain for the AURORAL project	Ontology Requirements	Ontology Repository	Ontology Issue Tracker	Ontology Releases	Biomass Payloads
AURORAL Logistic	This ontology aims to model the logistic domain for the AURORAL project	Ontology Requirements	Ontology Repository	Ontology Issue Tracker	Ontology Releases	Logistic Payloads
AURORAL Energy	This ontology aims to model the energy domain for the AURORAL project	Ontology Requirements	Ontology Repository	Ontology Issue Tracker	Ontology Releases	Energy Payloads
AURORAL Car-	This ontology aims to model the car	Ontology	Ontology Repository	Ontology Issue	Ontology	Car-booking

… including user-oriented artifacts

Turtle

TURTLE

```
@prefix ns1: <https://auroral.iot.linkeddata.es/def/core#> .
@prefix ns2: <https://www.w3.org/2019/wot/td#> .
@prefix ns3: <http://lexvo.org/ontology#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
ns1:Service1 a ns1:Service ;
    ns2:title "accomodationService"^^xsd:string ;
    ns2:description "service for accomodation"^^xsd:string ;
    ns2:created "12-01-2022T12:10:14"^^xsd:dateTime ;
    ns2:modified "12-01-2023T11:05:34"^^xsd:dateTime ;
    ns1:provider "Luxactive"^^xsd:string ;
    ns1:currentStatus "available"^^xsd:string ;
    ns1:hasDomain ns1:Tourism ;
    ns1:hasSubDomain "Musuemvisit"^^xsd:string ;
    ns1:hasRequirement "Photo not allowed"^^xsd:string ;
    ns1:serviceFree true ;
    ns1:numberOfDownload 15 ;
    ns1:versionOfservice "version1"^^xsd:string ;
    ns1:hasFunctionality "Guide in Musuem"^^xsd:string ;
    ns3:language   ns3:English ;
    ns3:language   ns3:Spanich ;
    ns1:applicableToGeographicalArea  ns1:PlaceService1 .

ns1:PlaceService1 a ns1:Place ;
    ns1:code "codeCountry"^^xsd:string ;
    ns1:name "nameCountry"^^xsd:string .
```

JSON-LD

JSON-LD

```
{
  "@context":"https://auroralh2020.github.io/auroral-ontology-contexts/core/services.json",
  "name":"Tourism monitor",
  "service_description":"Person counter",
  "provider":"BOSONIT",
  "status":"Available",
  "modified":"2021-11-09T18:25:43.511Z",
  "domain":"Mobility",
  "subdomain":"Fly",
  "functionalities":"Only read",
  "requirements":"The date to read the persons",
  "is_free":true,
  "url":"http://rur.tourism.com/itisveryimportant/birds",
  "downloaded": 1,
  "version":"1.4",
  "language":"spa",
  "place":"Spain"
}
```

AURORAL

u can find the list of interfaces involved in the data exchange for AURORAL project

Ontology	Exchanged Data	SHACL shapes	RDF examples (Turtle)	JSON-LD 1.1 examples	Context
AURORAL Core ontology	Data related to provide information about the auroral service.	Core Shape	Service in RDF	Service in Json-LD	Core context
AURORAL Privacy ontology	Data related to provide information about the privacy.	Privacy Shape	Privacy in RDF	Privacy in Json-LD	Privacy context
AURORAL Adapters	Data related to provide information about the Device and its measurement.	Adapters Shape	Adapters device in RDF	Adapters device in Json-LD	adapters context
AURORAL Tourism	Data related to provide information about the tourism activity.	Activity Shape	Tourism activity in RDF	Tourism Activity in Json-LD	Tourism context
AURORAL Marketplace ontology	Data related to provide information about the biomass marketplace.	marketplace Shape	Marketplace in RDF	Marketplace in Json-LD	Marketplace context
AURORAL Biomass	Data related to provide information about Biomass characteristics.	Biomass Shape	Biomass in RDF	Biomass in Json-LD	Biomass context
AURORAL Logistic ontology	Data related to provide information about logistic data.	Logistic Shape	RDF Example	Logistic in Json-LD	Logistic context
AURORAL Energy ontology	Data related to provide information about building energy consumption.	Energy Shape	Energy in RDF	Energy in Json-LD	Energy context
AURORAL car-booking	Data related to provide information about the car booking.	Car-booking Shape	Car-booking in RDF	Car-booking in Json-LD	Car-booking context

Context

```
{
  "@context" : {
    "core" : "https://auroral.iot.linkeddata.es/def/core#",
    "lexvo" : "http://lexvo.org/ontology#",
    "name" : {
      "@id": "core:serviceName",
      "@type": "xsd:string"
    },
    "service_description" : {
      "@id": "core:serviceDescription",
      "@type": "xsd:string"
    },
    "provider" : {
      "@id": "core:provider",
      "@type": "xsd:string"
    },
    "status" : {
      "@id": "core:currentStatus",
      "@type": "xsd:string"
    },
    "modified" : {
      "@id": "core:dateLastUpdate",
      "@type": "xsd:datetime"
    },
    "domain" : {
      "@id": "core:hasDomain",
      "@type": "xsd:string"
    },
    "subdomain" : {
      "@id": "core:hasSubDomain",
      "@type": "xsd:string"
    },
    "functionalities" : {
      "@id": "core:hasFunctionalities",
      "@type": "xsd:string"
    },
    "requirements" : {
      "@id": "core:hasRequirements",
      "@type": "xsd:string"
    },
    "is_free" : {
      "@id": "core:serviceFree",
      "@type": "xsd:boolean"
    },
```

SHACL

SHACL

```
@prefix :      <https://auroral.iot.linkeddata.es/def/core#> .
@prefix sh:    <http://www.w3.org/ns/shacl#> .
@prefix rdf:   <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix owl:   <http://www.w3.org/2002/07/owl#> .
@prefix xml:   <http://www.w3.org/XML/1998/namespace> .
@prefix afn:   <http://jena.apache.org/ARQ/function#> .
@prefix xsd:   <http://www.w3.org/2001/XMLSchema#> .
@prefix rdfs:  <http://www.w3.org/2000/01/rdf-schema#> .

<https://astrea.linkeddata.es/shapes#8f99085791371b49ddef1ccf7e0c33da>
        a             sh:PropertyShape ;
        sh:nodeKind  sh:IRIOrLiteral ;
        sh:path      :hasID .

<https://astrea.linkeddata.es/shapes#58bd5d9a299bb864bb0290d3337b267f>
        a             sh:PropertyShape ;
        sh:class     :Item ;
        sh:nodeKind  sh:IRIOrLiteral ;
        sh:path      :containsItem .

<https://astrea.linkeddata.es/shapes#a7465e1d9b24f8076950640e54abdd60>
        a             sh:PropertyShape ;
        sh:nodeKind  sh:IRIOrLiteral ;
        sh:path      :versionNumber .

<https://astrea.linkeddata.es/shapes#a7f69e63f216b0b7b5966485350e68d8>
        a                sh:NodeShape ;
        sh:nodeKind     sh:IRI ;
        sh:targetClass  :Municipality .

<https://astrea.linkeddata.es/shapes#b5198dd1f0944a5cc5366a1b45e3489e>
        a             sh:PropertyShape ;
        sh:nodeKind  sh:IRIOrLiteral ;
        sh:path      :hasLocation .

<https://astrea.linkeddata.es/shapes#6c6ec464a53850c4d3299239ca45f8bf>
        a             sh:PropertyShape ;
        rdfs:label   "has measurement" ;
        sh:class     <https://saref.etsi.org/core#Measurement> ;
        sh:name      "has measurement" ;
        sh:nodeKind  sh:BlankNodeOrIRI ;
        sh:path      <https://saref.etsi.org/core#hasMeasurement> .
```

Ontological resources support (if available) other research

CORAL: corpus of ontological requirements annotated with Lexico-Syntactic Patterns

- Includes:
  - Dictionary of 29 LSPs
  - A corpus of 834 requirements annotated with LSPs
- Openly available in HTML, CSV and RDF
- Already applied in Themis for testing ontologies

CORAL

http://coralcorpus.linkeddata.es/

Characteristic	Example
Identifier of the LSP	LSP-SC-EN
Description	The definition of a subsumption relation in an ontology
NeOn ontology design pattern identifier	"LP-SC-01", which represent the ODP related to subclassOf relations
BNF formalization	**There are QUAN CN-CATV NP<superclass> PARA [(NP<subclass>,)* and] NP<subclass>"**
Examples	There are different types of devices: sensor and actuator
OWL Constructs	**subClassOf, Class (Thing, Nothing)**
DL Expressivity	AL

Fernández-Izquierdo A.; Poveda-Villalón M.; García-Castro R. CORAL: A Corpus of Ontological Requirements Annotated with Lexico-Syntactic Patterns. 16th Extended Semantic Web Conference (ESWC 2019). Resources track, LNCS 11503. Portoroz, Slovenia. June 2019.

Opening ontology development to others

Ontology engineers

End users

Domain experts

Software engineers

...

INSIDERS

OUTSIDERS

Local/regional government

Public/private/ civic sectors

Citizens

Standardization

...

SAREF4CITY

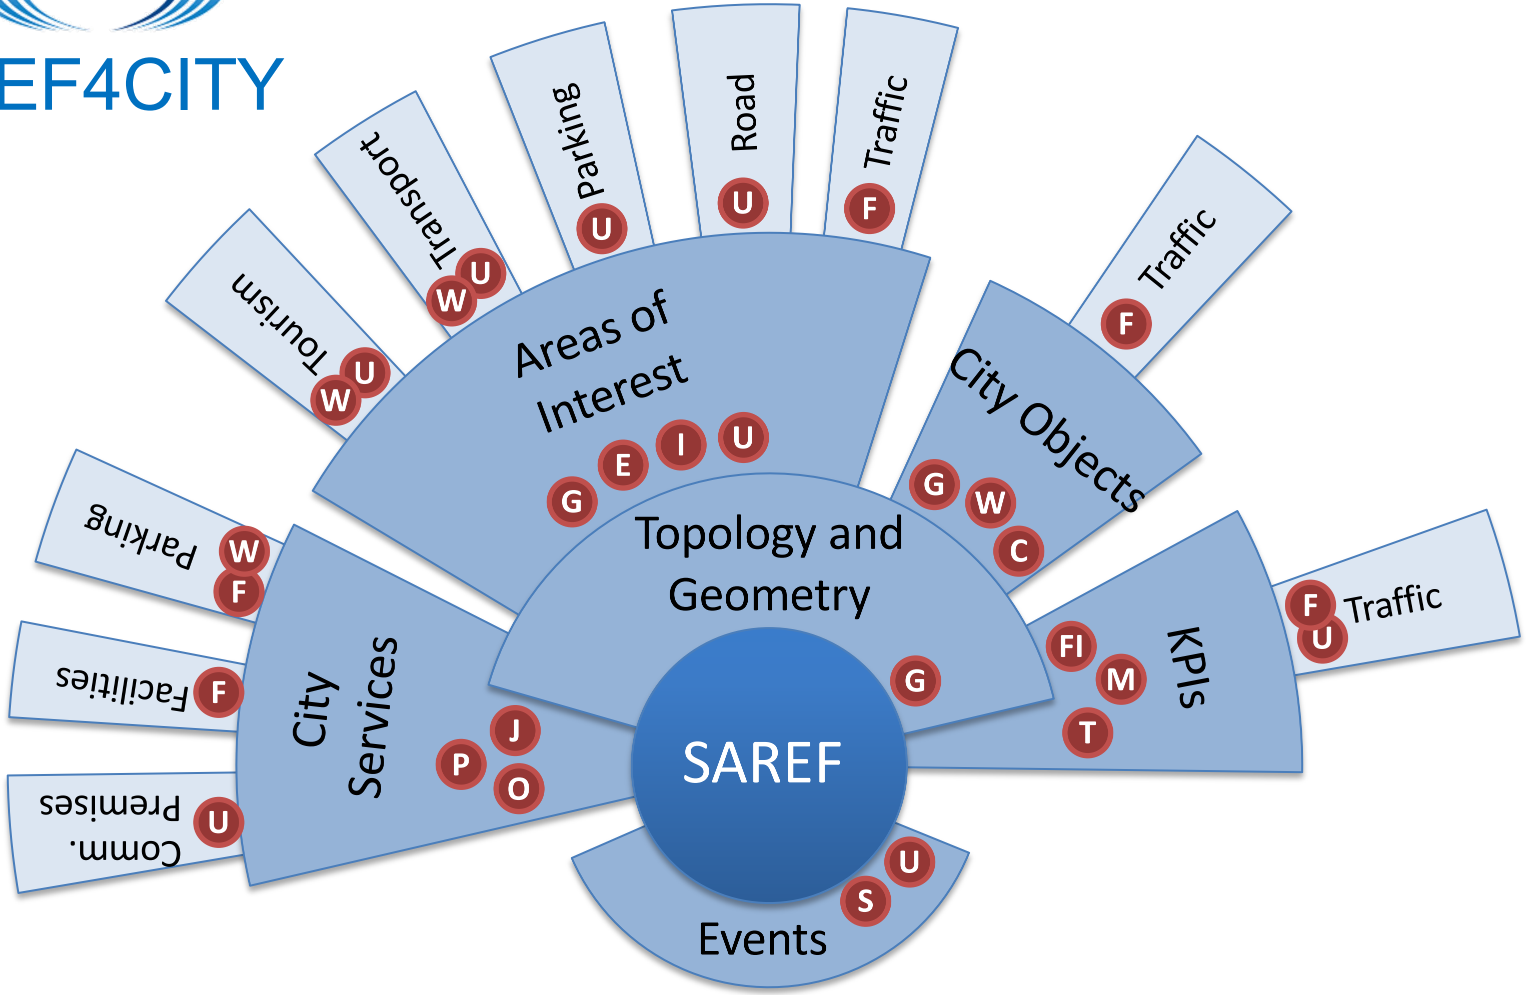

Requirements:

- E EU Metadata Registry
- F FEMP Open Data Guide exemplary datasets
- FI FIWARE data model for KPIs
- I ISA Programme Location Core Vocabulary
- J Joinup Core Public Organization Vocabulary
- P Joinup Core Public Service Vocabulary
- C OGC CityGML
- G OGC GeoSPARQL
- S schema.org
- U Vocabulary referenced by AENOR UNE 178301:2015
- O W3C Registered Organization Vocabulary
- W W3C WGS84 Geo Positioning vocabulary
- M ISO/IEC 30182:2017
- T ITU-T Y.4903/L.1603 (10/2016)

Espinoza-Arias P, Poveda-Villalón M, García-Castro R, Corcho O. Ontological Representation of Smart City Data: From Devices to Cities. Applied Sciences. 2019; 9(1):32. https://doi.org/10.3390/app9010032

Ontological commitments

- Agreements to use the vocabulary in a coherent and consistent manner
- Connection between the ontology vocabulary and the meaning of the terms of such vocabulary
- An agent commits (conforms) to an ontology if it "acts" consistently with the definitions

How do we represent and validate them?

- A priori + a posteriori

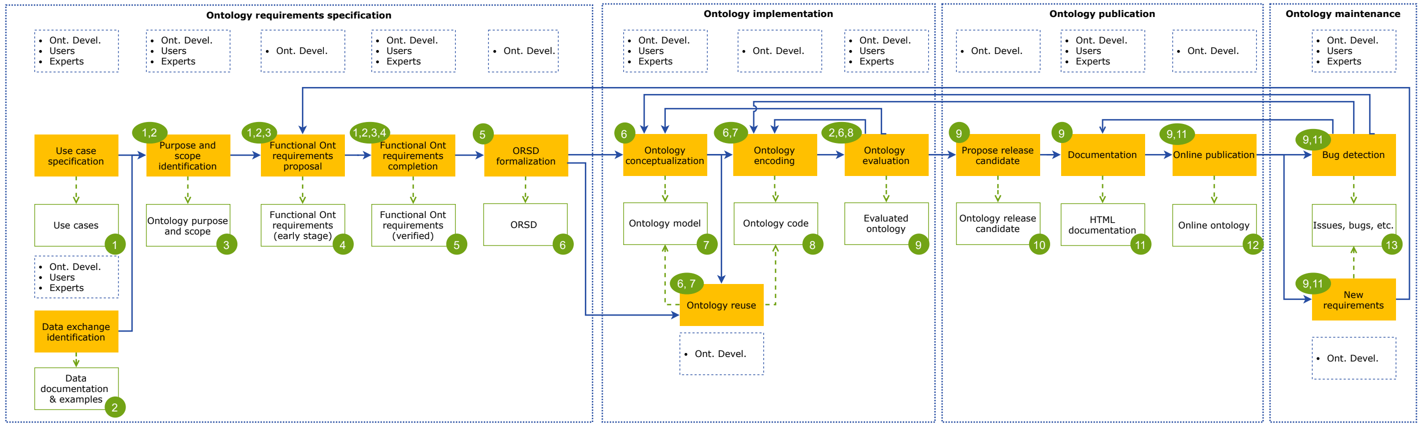

Gruber, T.; Olsen, G. An Ontology for Engineering Mathematics. Fourth International Conference on Principles of Knowledge Representation and Reasoning. Ed by Doyle and Torasso. Morgan Kaufmann. 1994. Also as KSL-94-18.
Guarino, N.; Carrara, M.; Giaretta, P. Formalizing Ontological Commitments. 12th National Conference on Artificial Intelligence. AAAI-94. 1994. 560-567

- **What is the difference between what is explicitly defined in the requirements and the consequent ontology implementation?**
  - Rate of related OWL constructs in **requirements**: **15-35%**
    - The analysed requirements were vague and cover a small set of OWL constructs
  - Rate of related OWL constructs in the **ontology**: **38-69%**
    - Ontology engineers take a considerable amount of modelling decisions that are not deduced from the requirements

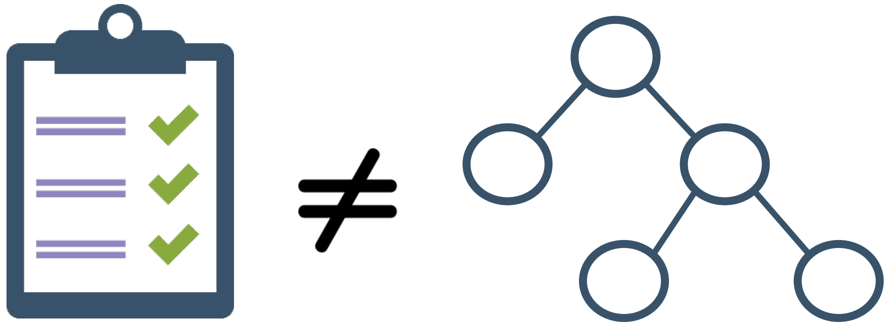

Fernández-Izquierdo A.; Poveda-Villalón M.; García-Castro R. Analysing Ontological Requirements: A Journey from Requirements to Code and Back. XIX Conference of the Spanish Association for Artificial Intelligence (CAEPIA 20/21). ISBN: 978-84-09-30514-8. Málaga, Spain. Septiembre 2021.

Need ontology verification

- http://oops.linkeddata.es/
- Implements (semi)automatic detection of 33 pitfalls (of 41)
- Available as:
  - Web app + REST API

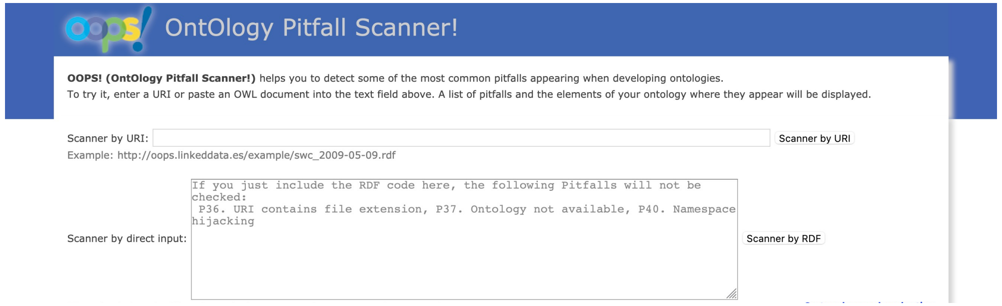

Poveda-Villalón, María, Asunción Gómez-Pérez, and Mari Carmen Suárez-Figueroa. "Oops!(ontology pitfall scanner!): An on-line tool for ontology evaluation." Int. Journal on Semantic Web and Information Systems 10.2 (2014): 7-34.

- http://themis.linkeddata.es/
- Automates test implementation and execution activities
- Available as:
  - Web app + REST API + CLI

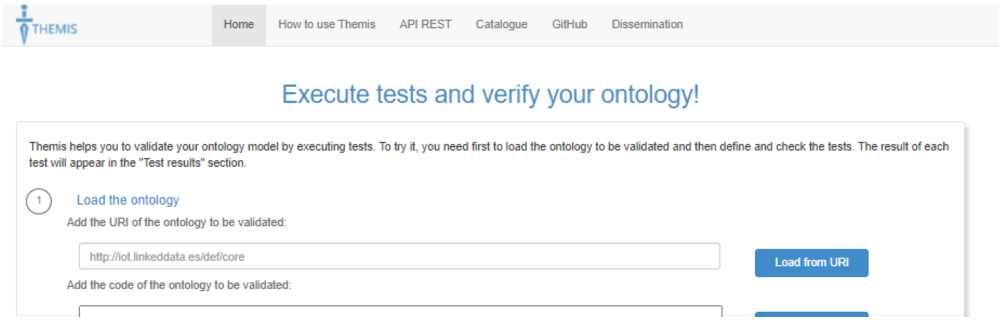

Alba Fernández-Izquierdo, Raúl García-Castro, Conformance testing of ontologies through ontology requirements, Engineering Applications of Artificial Intelligence, Volume 97, 2021, ISSN 0952-1976.

Need ontology validation on many fronts

▪ Validating requirements

Users without any ontology background find it difficult to understand ontology restrictions

Testing frameworks reduce the number of errors during the verification process for users familiar and experts in OWL

Everyone (even OWL experts!) makes mistakes

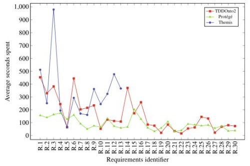
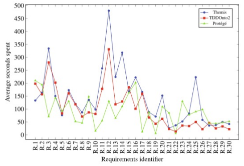
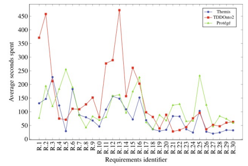

A. Fernández-Izquierdo, R. García-Castro. Ontology verification testing using lexico-syntactic patterns, Information Sciences, Vol. 582, 2022, pp. 89-113, ISSN 0020-0255.

▪ Validating conceptualization

Chowlk

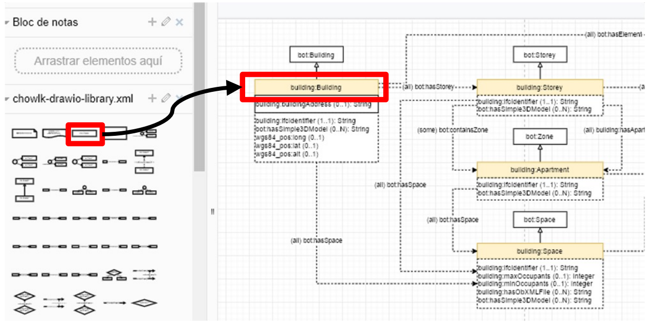

https://chowlk.linkeddata.es/

Chávez-Feria, S., García-Castro, R., Poveda-Villalón, M. (2022). Chowlk: from UML-Based Ontology Conceptualizations to OWL. In: , et al. The Semantic Web. ESWC 2022. Lecture Notes in Computer Science, vol 13261. Springer.

▪ Validating conformance

Expressive axioms/requirements are not shared

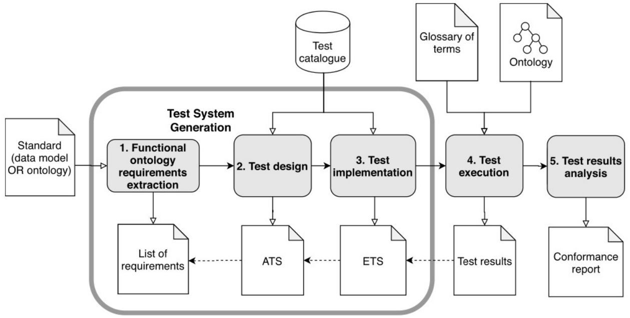

Alba Fernández-Izquierdo, Raúl García-Castro. Conformance testing of ontologies through ontology requirements, Engineering Applications of Artificial Intelligence, Vol. 97, 2021.

▪ Validating data exchanges

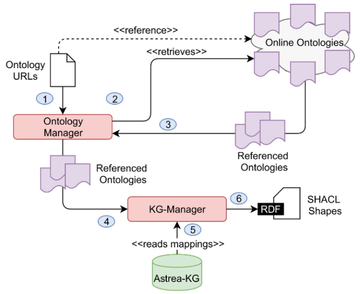

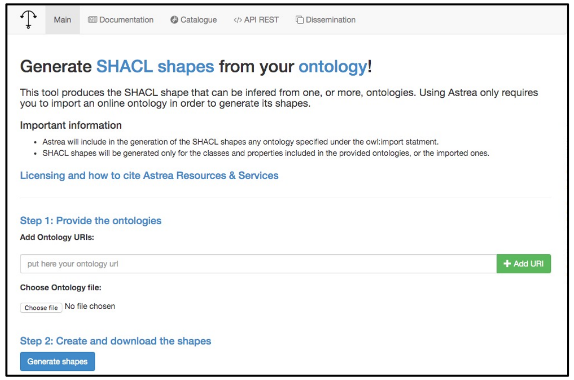

http://astrea.linkeddata.es/

Cimmino A.; Fernández-Izquierdo A.; García-Castro R. Astrea: automatic generation of SHACL shapes from ontologies. 17th Extended Semantic Web Conference (ESWC 2020). Resources track, LNCS 12123. Springer. Heraklion, Greece. June 2020.

Community-driven ontology engineering is a long-distance race

- Almost 10 years since the publication of SAREF v1
- Lessons learnt from the development of the SAREF ontologies
  - Reduced scope of the extensions: Fast development → fast evolution
  - Strategy for ontology modularization not trivial
  - Implicit Ontology Design Patterns could be formalised
  - Stakeholder workshops better face to face (also create community)
  - Tool support and training are essential
  - and many more …

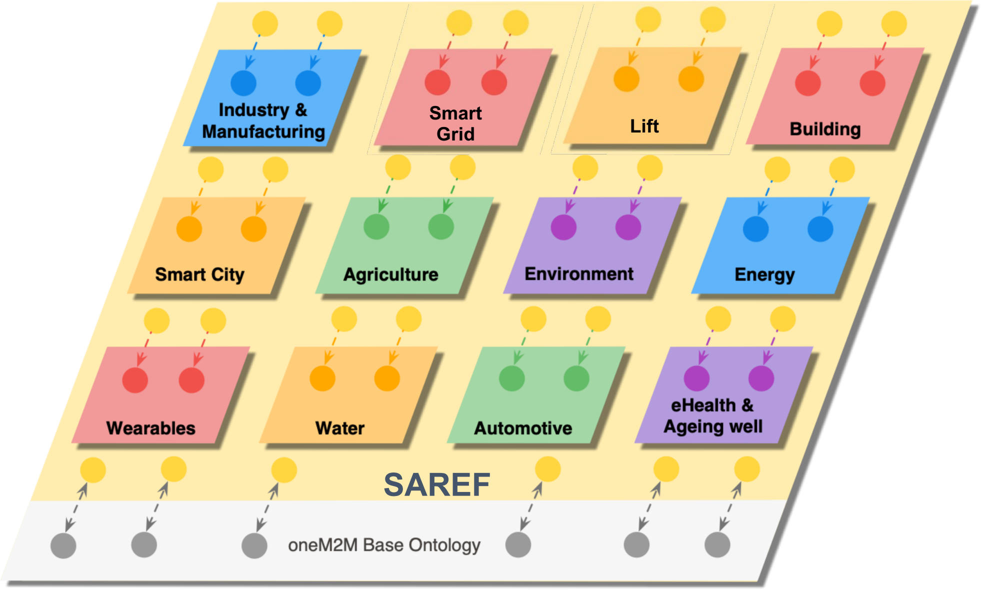

García Castro R.; Lefrançois M.; Poveda-Villalón M.; Daniele L. The ETSI SAREF Ontology for Smart Applications: A Long Path of Development and Evolution. Energy Smart Appliances: Applications, Methodologies, and Challenges. Wiley, pp. 183-216. June 2023

Towards ontology governance

SAREF governance

SPECIFICATIONS

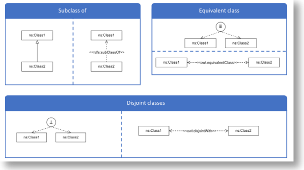

Principles

- Standards
- Requirements
- Guidelines

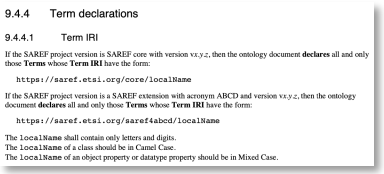

9.4.4 Term declarations

9.4.4.1 Term IRI

If the SAREF project version is SAREF core with version v.x.y.z, then the ontology document **declares** all and only those **Terms** whose **Term IRI** have the form:

`https://saref.etsi.org/core/localName`

If the SAREF project version is a SAREF extension with acronym ABCD and version v.x.y.z, then the ontology document **declares** all and only those **Terms** whose **Term IRI** have the form:

`https://saref.etsi.org/saref4abcd/localName`

The `localName` shall contain only letters and digits.
The `localName` of a class should be in Camel Case.
The `localName` of an object property or datatype property should be in Mixed Case.

Processes

- Actors
- Ontology development methodology
- Workflows

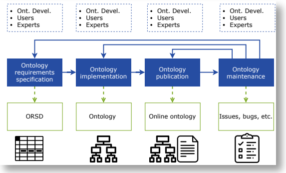

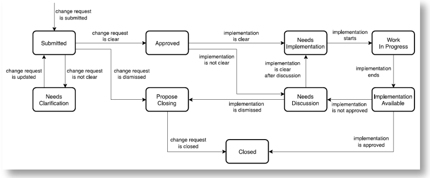

Technology

- Development framework
- Technical requirements

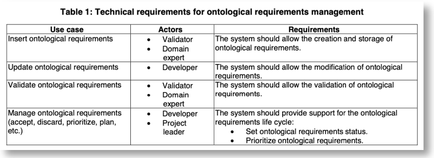

Table 1: Technical requirements for ontological requirements management

Use case	Actors	Requirements
Insert ontological requirements	• Validator • Domain expert	The system should allow the creation and storage of ontological requirements.
Update ontological requirements	• Developer	The system should allow the modification of ontological requirements.
Validate ontological requirements	• Validator • Domain expert	The system should allow the validation of ontological requirements.
Manage ontological requirements (accept, discard, prioritize, plan, etc.)	• Developer • Project leader	The system should provide support for the ontological requirements life cycle: • Set ontological requirements status. • Prioritize ontological requirements.

ETSI TS 103 411

ETSI TS 103 608

ETSI TS 103 673

SOFTWARE

SAREF forge

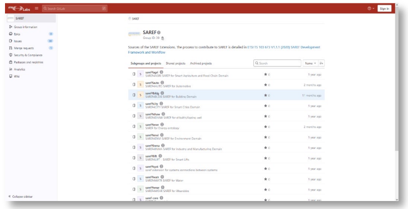

https://labs.etsi.org/rep/saref/

SAREF ontology portal

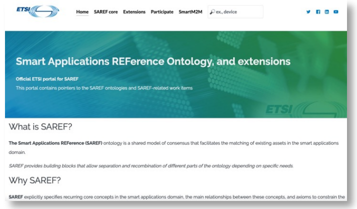

https://saref.etsi.org/

SAREF pipeline

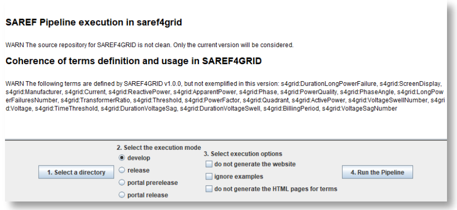

THEMIS

https://labs.etsi.org/rep/saref/saref-pipeline/

Semantic **interoperability**

IEEE Standard Computer Dictionary

Interoperability

The ability of two or more systems or components to exchange information and to use the information that has been exchanged.

IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries (New York, NY: 1990)

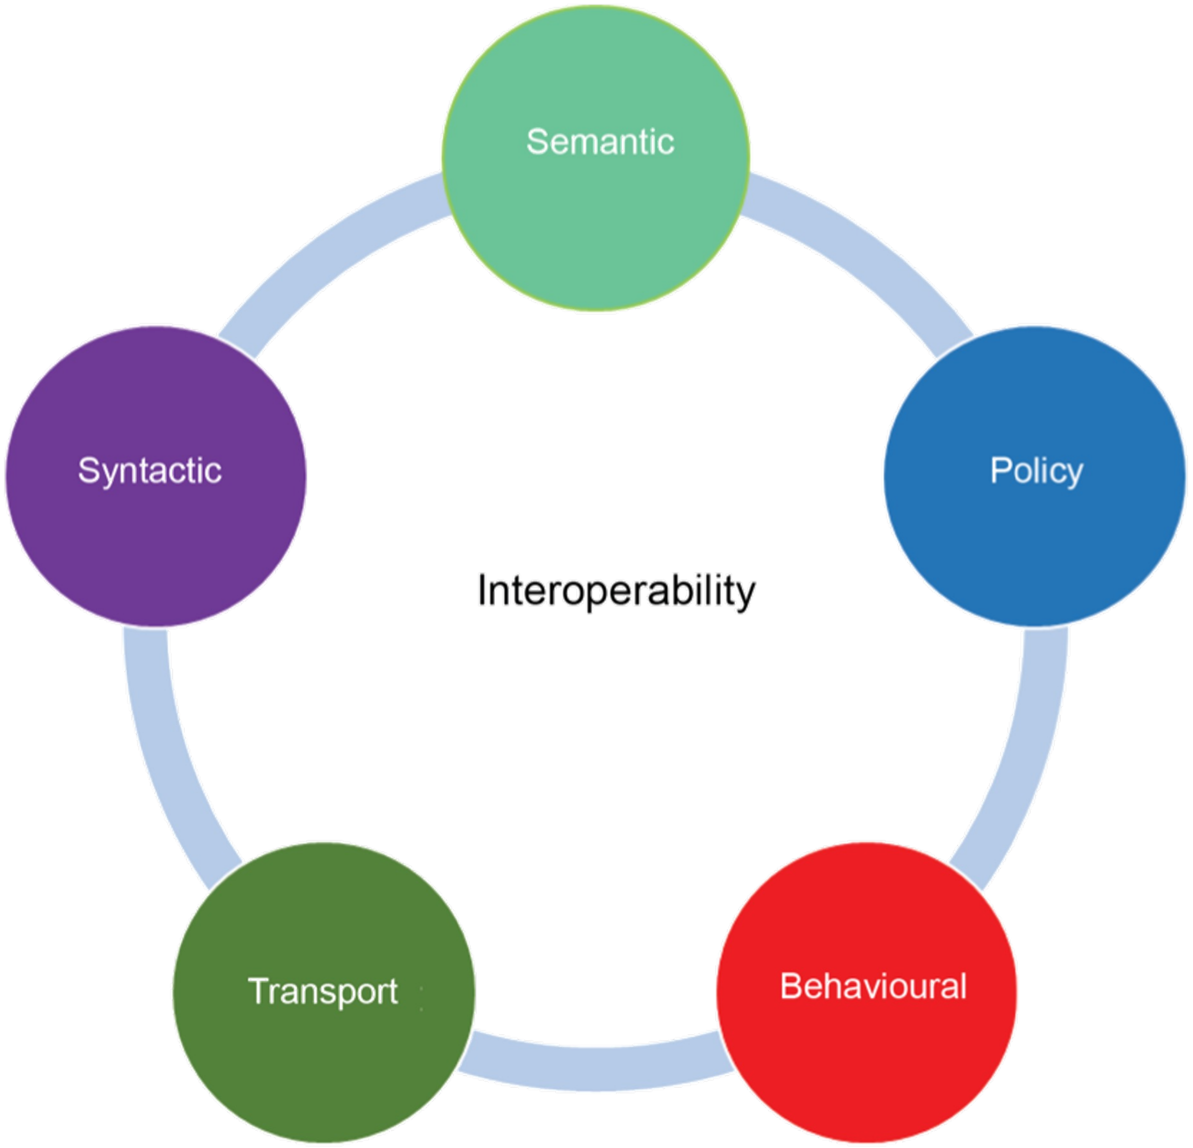

ISO/IEC 21823-1:2019 Internet of things (IoT) — Interoperability for IoT systems — Part 1: Framework

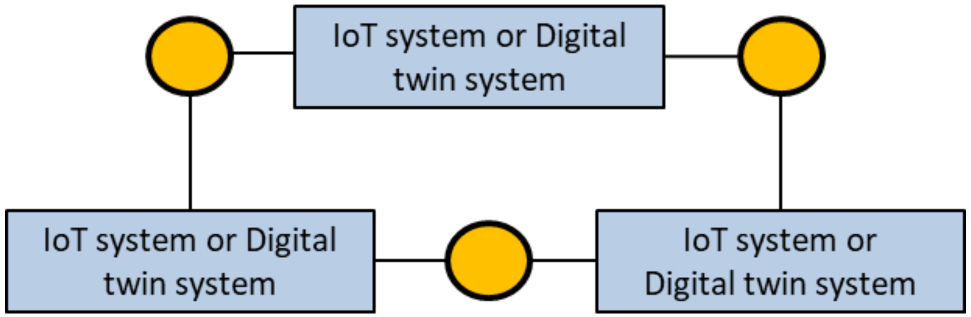

System of systems scenario

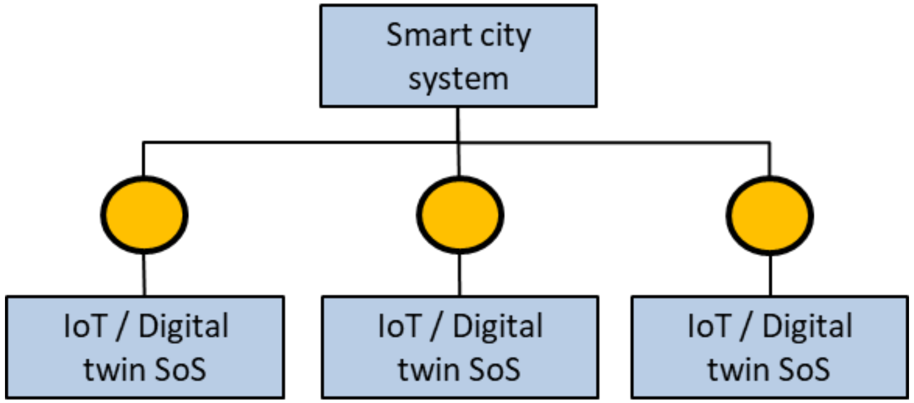

Smart city scenario

ISO/IEC 21823-5 Internet of things – Interoperability for IoT systems – Part 5: Behavioural and policy interoperability. Editors: A. Kung, J. Alve, K. Tobich, J. Lee, R. García-Castro, J. deMeer

WORK IN PROGRESS

Need shared interoperability specifications: interoperability profiles

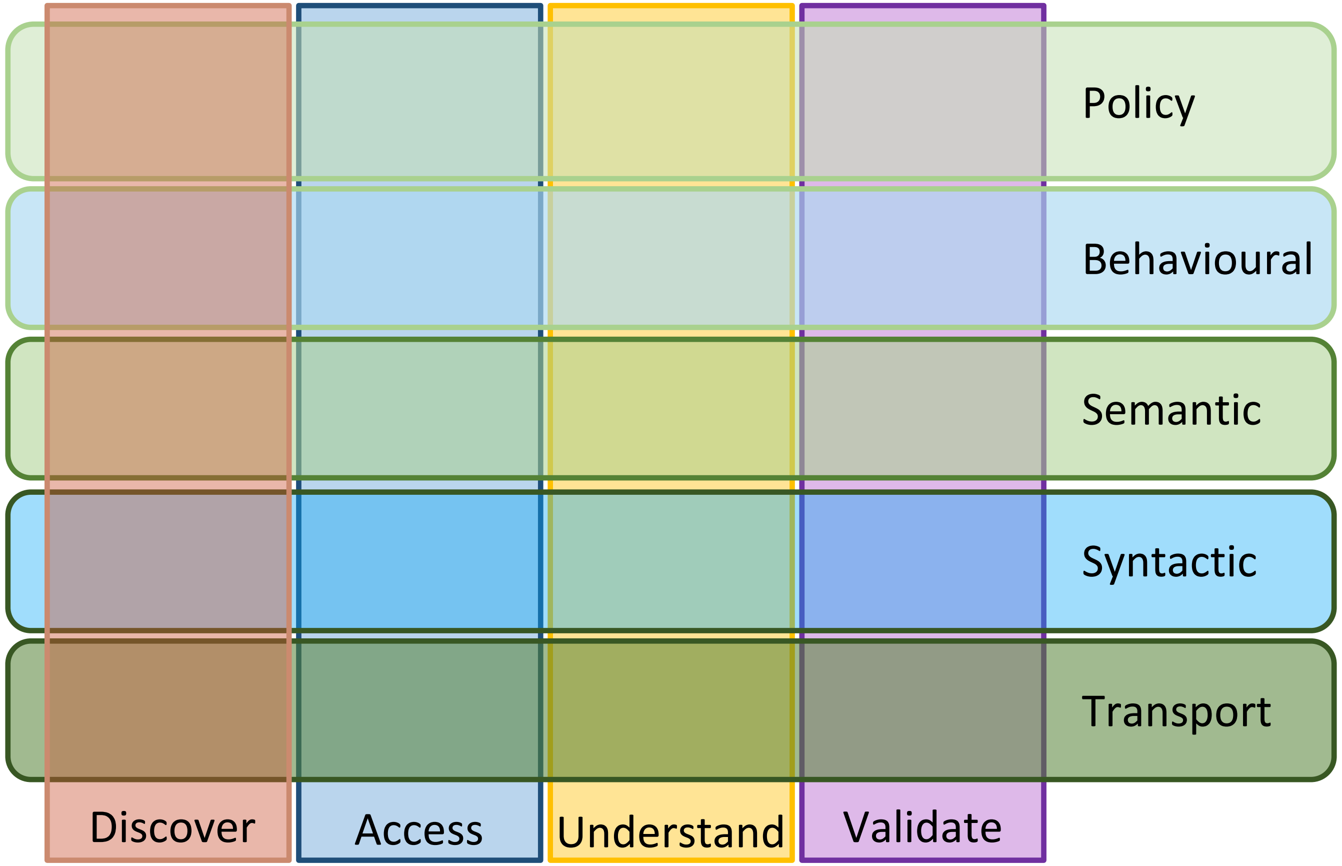

AURORAL interoperability profile

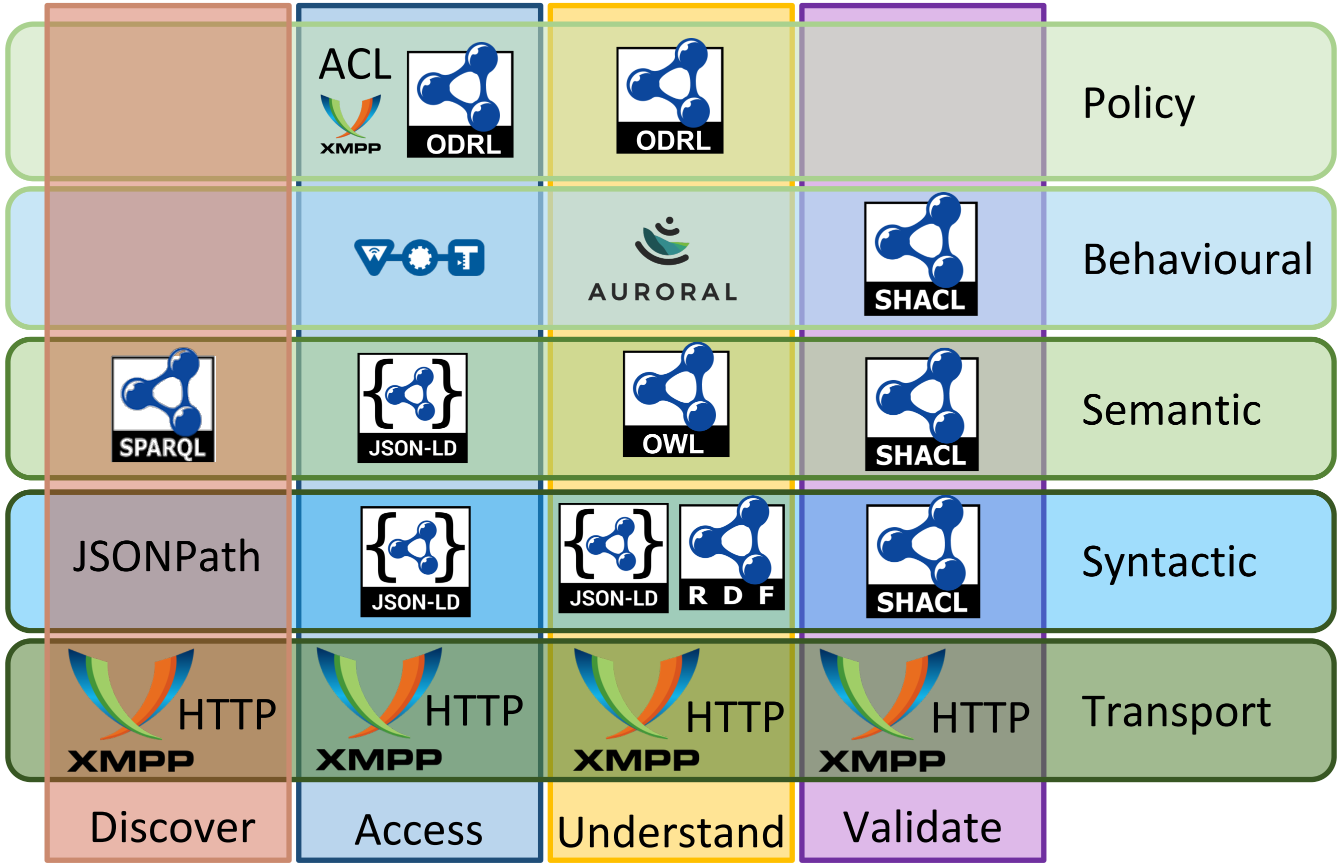

Need for semantic discovery

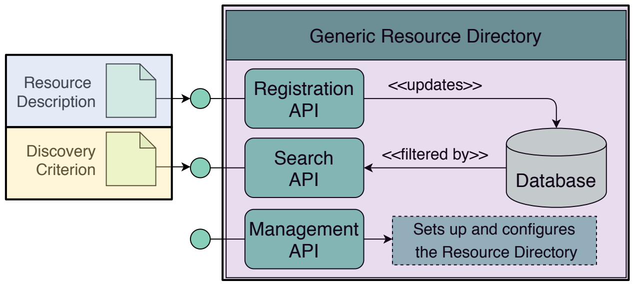

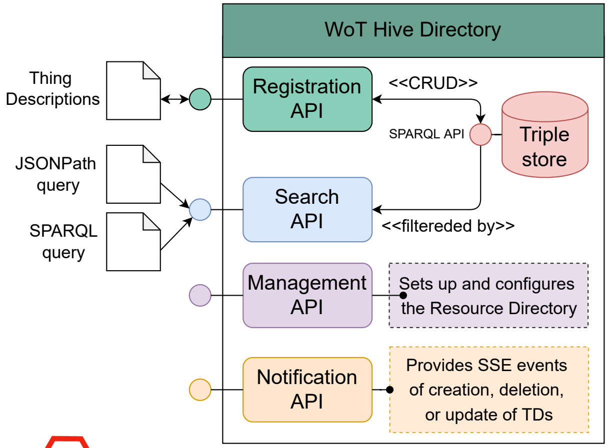

https://github.com/oeg-upm/wot-hive

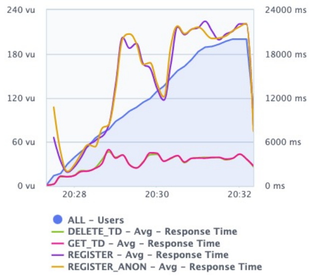

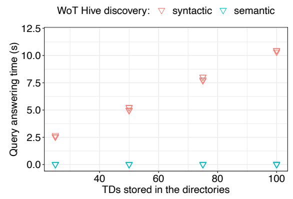

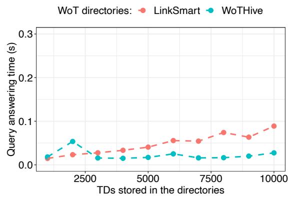

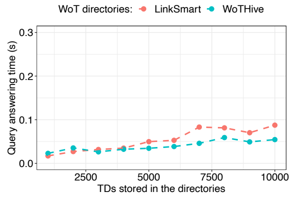

- Why semantic Web of Things discovery?
  - Based on a standard language and protocol (SPARQL)
  - Allows to express more complex queries (filtering/JSON Path vs querying/SPARQL)
  - Query federation for decentralising discovery (using standards, SPARQL)
- Is it feasible?
  - Semantic discovery seems to outperform syntactic one in
    - Complexity of discovery criteria (queries) that can be expressed
    - Query answering time
  - However, when the query answer has a large size syntactic discovery is faster
    - Due to the verbosity of the answer

Cimmino A., García-Castro R. WoTHive: Enabling Syntactic and Semantic Discovery in the Web of Things. Open Journal of Internet of Things. Vol. 8 (1), pp. 54-65. RonPub. September 2022

Can we decentralise discovery?

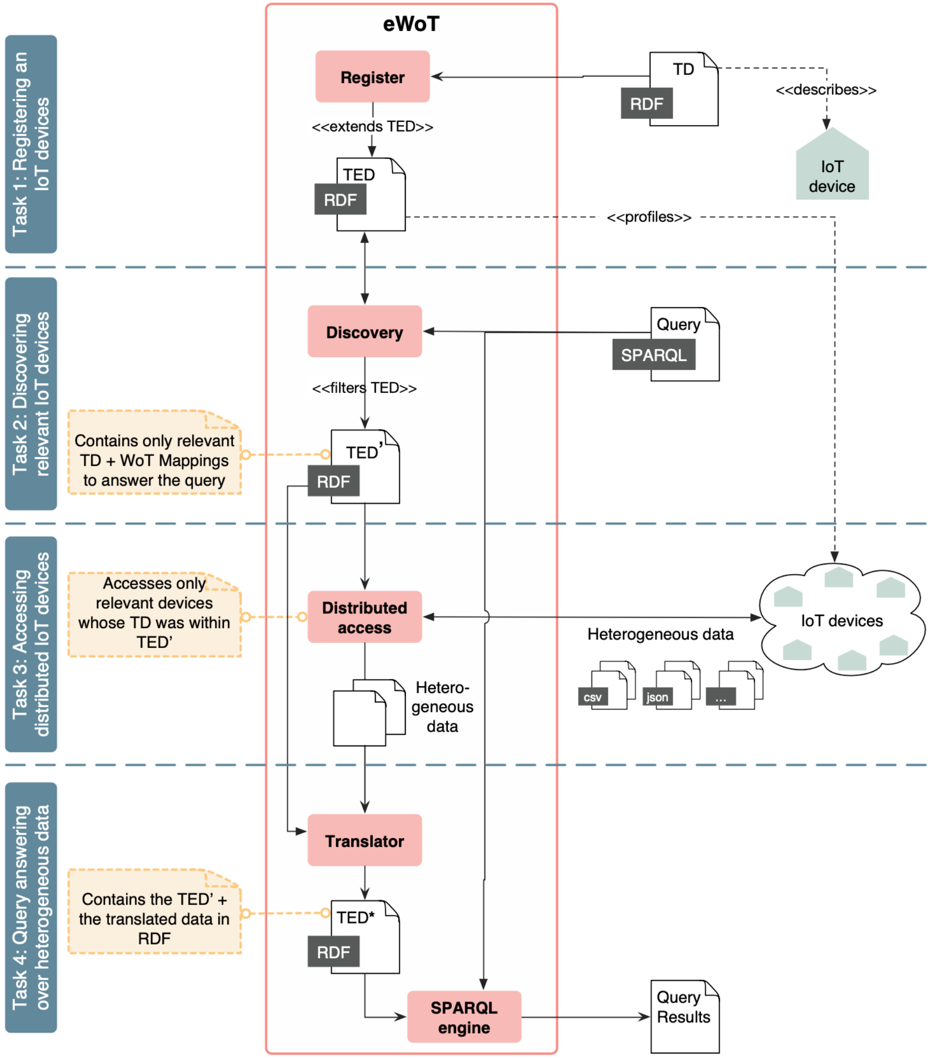

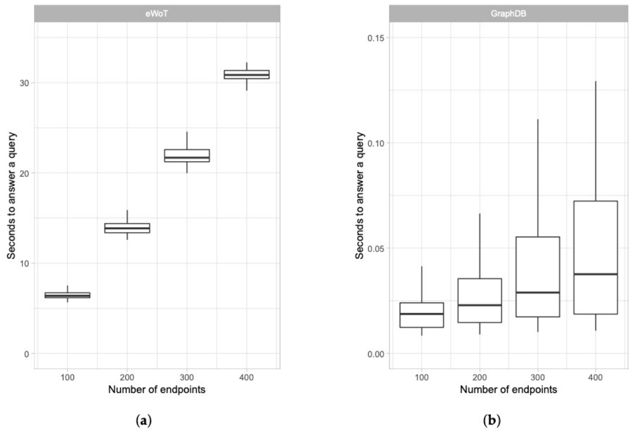

Figure 6. Comparison of eWoT and centralised approach based on GraphDB: (**a**) eWoT query-answering time and (**b**) GraphDB query answering time.

Table 3. Percentage of Discovery time and Distributed Access plus Translation time.

	Query Answering Time in %									
	100 Endpoints		250 Endpoints		500 Endpoints		750 Endpoints		1000 Endpoints	
Query Type	Discovery	Access	Discovery	Access	Discovery	Access	Discovery	Access	Discovery	Access
Linear 1	95%	5%	96%	4%	96%	4%	96%	4%	95%	5%
Linear 2	96%	4%	96%	4%	96%	4%	96%	4%	96%	4%
Linear 3	95%	5%	96%	4%	96%	4%	96%	4%	96%	4%
Linear 4	95%	5%	96%	4%	96%	4%	96%	4%	96%	4%
Star 1	96%	4%	96%	4%	96%	4%	96%	4%	96%	4%
Star 2	96%	4%	96%	4%	96%	4%	96%	4%	96%	4%
Star 3	96%	4%	96%	4%	96%	4%	96%	4%	96%	4%
Star 4	96%	4%	96%	4%	96%	4%	96%	4%	96%	4%
Tree 1	96%	4%	96%	4%	96%	4%	96%	4%	96%	4%
Tree 2	95%	5%	96%	4%	96%	4%	96%	4%	96%	4%
Tree 3	96%	4%	96%	4%	96%	4%	96%	4%	96%	4%
Tree 4	96%	4%	96%	4%	96%	4%	96%	4%	96%	4%
Complex 1	93%	7%	94%	6%	96%	4%	95%	5%	96%	4%
Complex 2	95%	5%	96%	4%	96%	4%	96%	4%	96%	4%
Complex 3	95%	5%	96%	4%	96%	4%	96%	4%	96%	4%
Complex 4	95%	5%	96%	4%	96%	4%	96%	4%	96%	4%
Complex 5	95%	5%	96%	4%	96%	4%	96%	4%	96%	4%
Complex 6	95%	5%	96%	4%	96%	4%	96%	4%	96%	4%
Complex 7	96%	4%	96%	4%	96%	4%	96%	4%	96%	4%
Complex 8	95%	5%	96%	4%	96%	4%	96%	4%	96%	4%

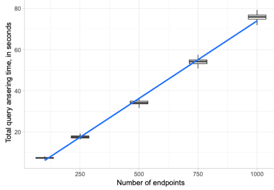

Figure 5. Averaged results: whiskers and plot of the results obtained handling different numbers of endpoints.

- Centralised discovery
  - Faster compared to decentralised (discovery+access+translate+query)
- Decentralised discovery
  - Takes time; access is nearly instant
  - Stable query answering time and scales linearly
  - Ensures data freshness (not guaranteed by centralised)
  - Easier integration of new devices

Cimmino A.; Poveda-Villalón M.; García-Castro R. eWoT: A Semantic Interoperability Approach for Heterogeneous IoT Ecosystems Based on the Web of Things. Sensors. Vol. 20(3), pp. 1-19. MDPI. February 2020

Exposing linked data for interoperability

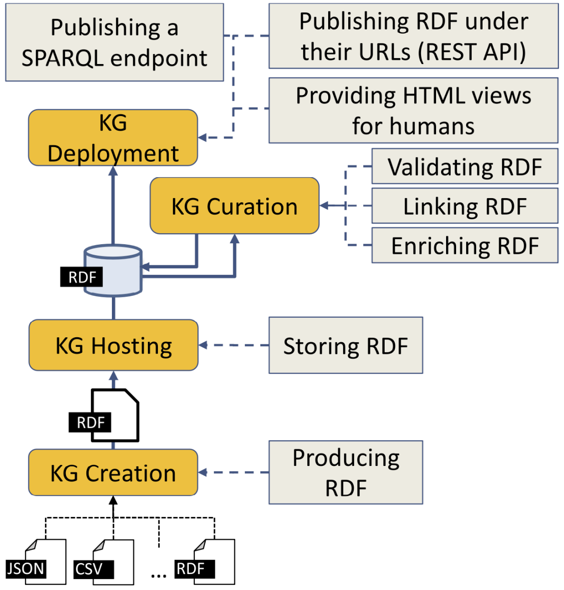

Fig. 1. KG life cycle [34] and related tasks.

Table 1

Elicited requirements met by existing tools types

KG life cycle	Categories of tools from the literature					
	Requirements	RDF materialisers	OBDI/A	RDF frameworks	RDF triple stores	RDF publishers
KC Creation	R01	~	–	~	–	–
	R02	✓	–	~	–	–
	R03	~	~	~	–	–
	R04	~	~	–	–	–
	R05	~	–	–	–	–
	R06	–	–	–	–	–
	R07	–	–	–	–	–
KG Hosting	R08	–	–	~	–	–
	R09	–	–	–	–	–
KG Curation	R10	–	–	~	~	–
KG Deployment	R11	–	–	~	~	–
	R12	–	~	~	✓	~
	R13	–	–	~	✓	~
	R14	–	–	–	–	~
	R15	–	–	–	–	~
	R16	–	–	–	–	–

- KG systems still lack some desirable features:
  - Support the whole knowledge graph life cycle
  - Support multiple mapping languages
  - Bi-directional (read-write) mappings
  - Mappings enhanced with programming languages
  - Extensibility: querying, linking, validating, policies, etc.

https://github.com/helio-ecosystem

Cimmino A.; García-Castro R. Helio: a framework for implementing the life cycle of knowledge graphs. Semantic Web – Interoperability, Usability, Applicability. Vol. 15 (1), pp. 223-249. IOS Press. January 2024

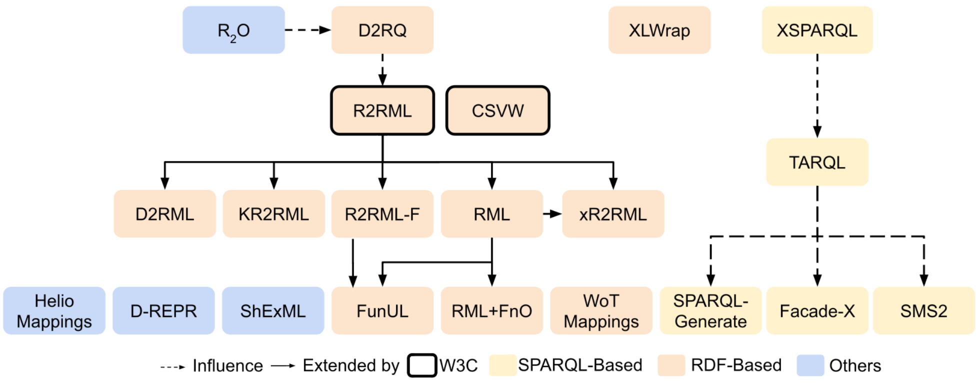

Fig. 1. Existing mapping languages and their relationships.

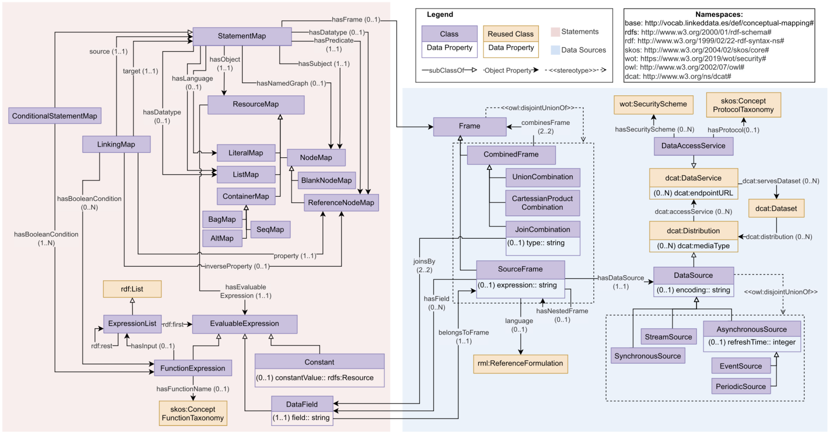

http://vocab.linkeddata.es/def/conceptual-mapping

- Do declarative mapping languages share common inherent characteristics?
  - Up to some extent, yes
- Can they be modelled?
  - The Conceptual Mapping ontology aims to represent the expressiveness of mapping languages
- What is this useful for?
  - Mapping translation and interoperability (requires formal semantics and operators)
  - Enhancement of knowledge graph construction workflows
  - Support mapping sharing and management (e.g., MappingPedia)

Iglesias-Molina A.; Cimmino A.; Ruckhaus E.; Chaves-Fraga D.; García-Castro R.; Corcho O. An Ontological Approach for Representing Declarative Mapping Languages. Semantic Web – Interoperability, Usability, Applicability. Vol. 15 (1) pp. 191-221. January 2024

- ODRL enables descriptive data usage policies
  - Information model: https://www.w3.org/TR/odrl-model/
  - Vocabulary: https://www.w3.org/TR/odrl-vocab/
- Limitations of ODRL:
  - Lacks expressiveness for injecting dynamic external data in policies
  - No implementation specification for constraints
  - No details on policy evaluation and enforcement

Policy interoperability is still open research

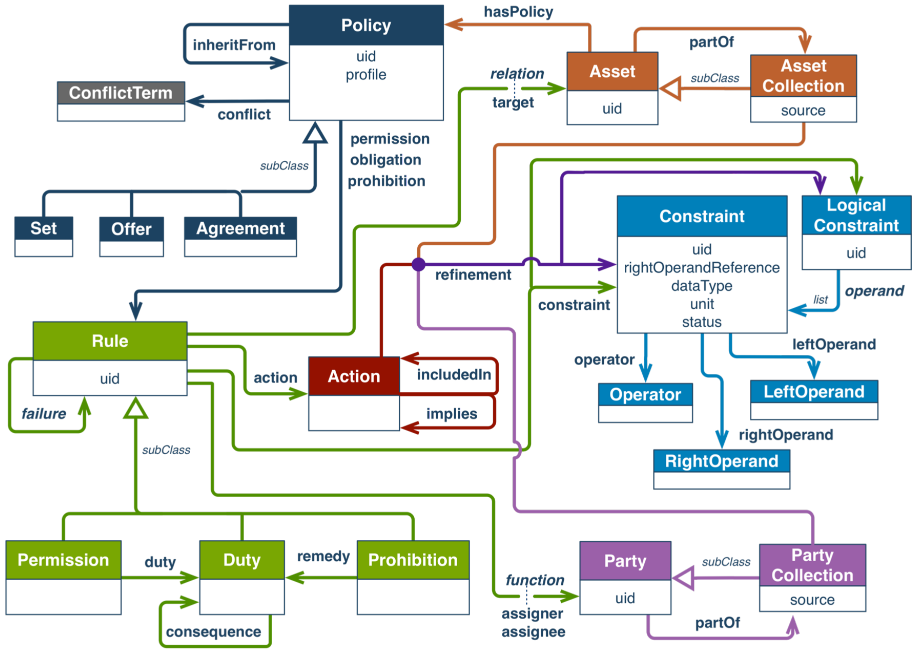

Cimmino A.; Cano-Benito J.; García-Castro R. Practical challenges of ODRL and potential courses of action. International Workshop on Trusting Decentralized Knowledge Graphs and Web Data (TrusDeKW 2023). WWW '23 Companion: Companion Proceedings of the ACM Web Conference 2023, pp. 1428-1431. ACM. Austin, TX, USA. 1 May 2023

Dynamic and decentralised privacy enforcement

AURORAL

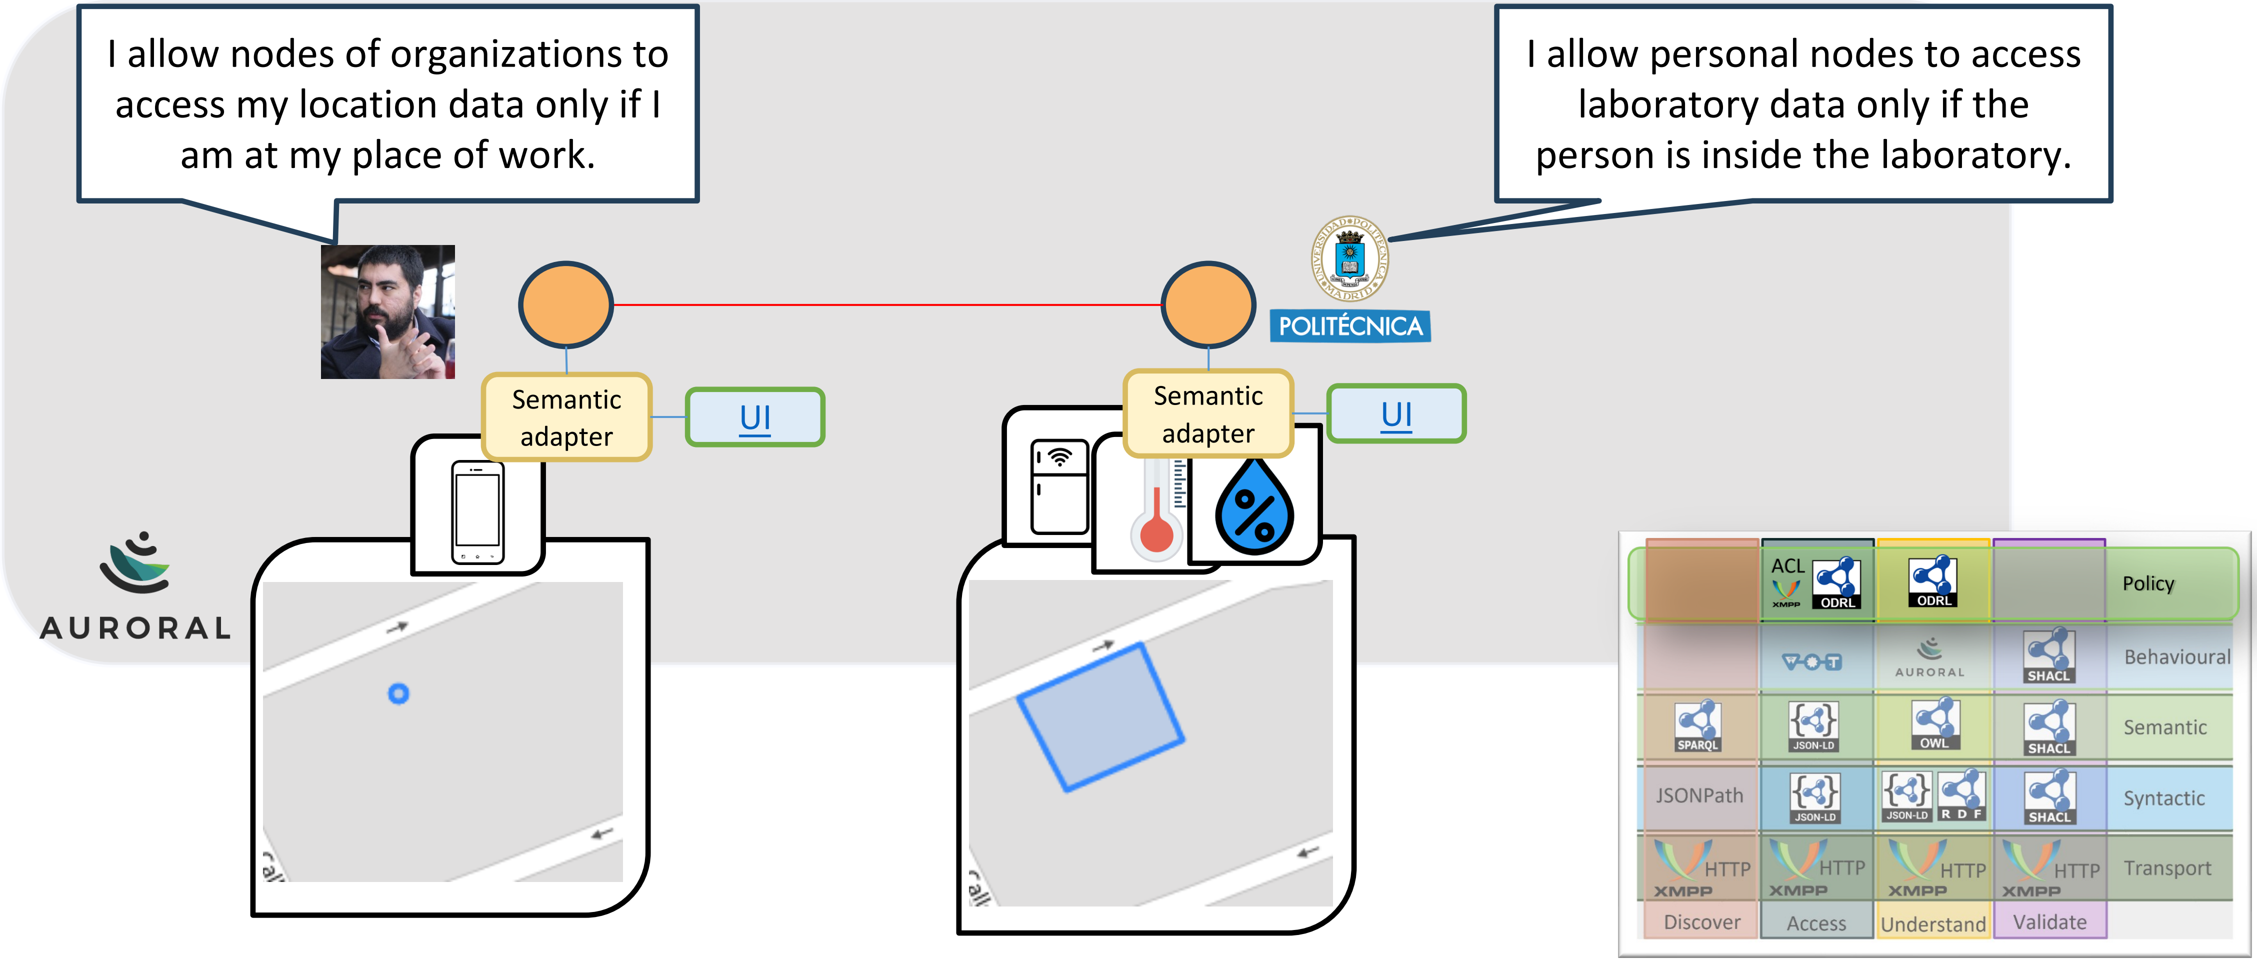

https://github.com/ODRE-Framework

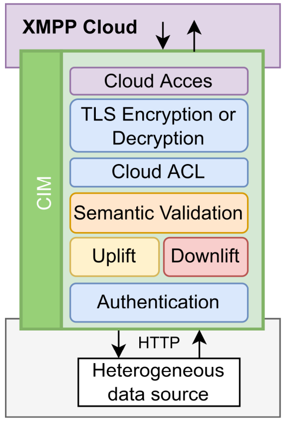

Fig. 1. CIM architecture.

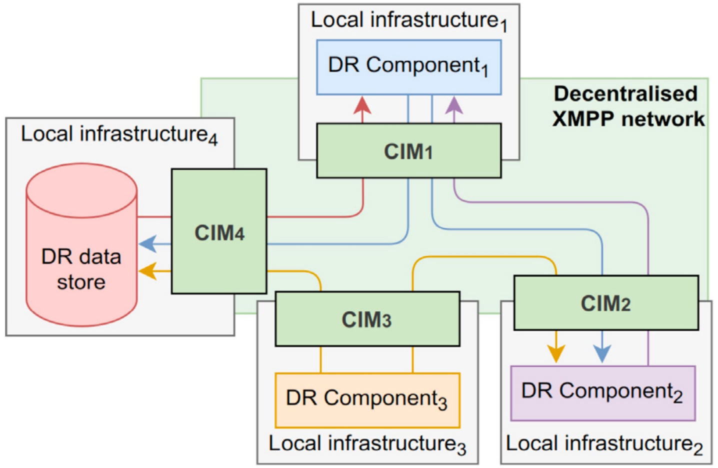

Fig. 4. A decentralised DR system using the CIM.

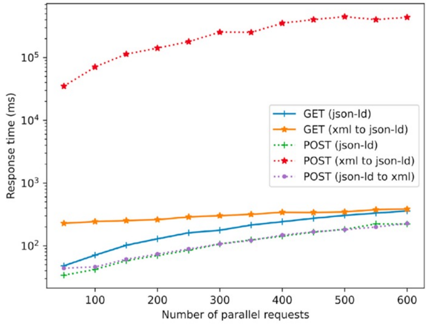

Fig. 7. Response times for scaling parallel requests.

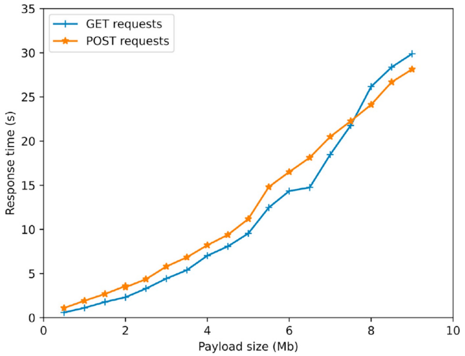

Fig. 8. Response times for scaling a payload size.

- Decoupled semantic interoperability modules:
  - Seems to be a suitable approach
  - In charge of uplifting, downlifting, validation…
  - Reusable up to some extent (vs adapters that are non-reusable)
- Developers tend to develop ad-hoc uplifting solutions and usually do not consider downlifting
  - Requires more dissemination and training
- Acknowledge the limitations of XMPP networks:
  - In high-latency scenarios the time for data exchanges was not low enough
  - Not suitable for large payloads (e.g., historical data)

Cimmino A.; Cano-Benito J.; Fernández-Izquierdo A.; Patsonakis C.; Tsolakis A.; García-Castro R.; Ioannidis D.; Tzovaras D. A scalable, secure, and semantically interoperable client for cloud-enabled Demand Response. Future Generation Computer Systems. Vol. 141, pp. 54-66. April 2023.

Scenarios to combine semantic technologies with blockchain

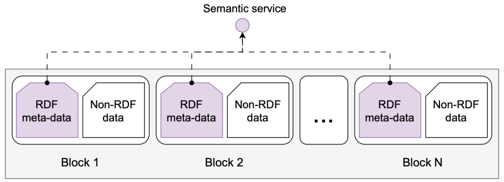

Figure 4.1: Blockchain with semantic metadata

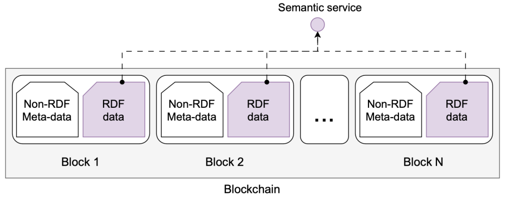

Figure 4.2: Blockchain with semantic data

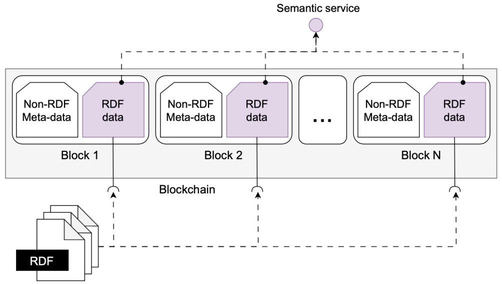

Figure 4.3: Blockchain with semantic data into the block directly

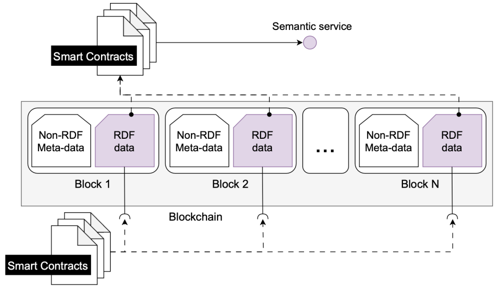

Figure 4.4: Blockchain with semantic data into the block using semantic smart contracts

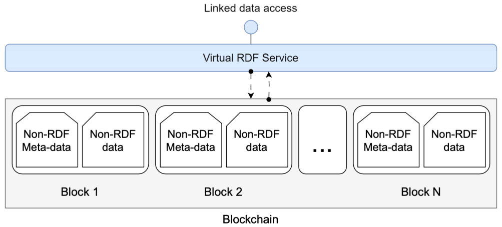

Figure 4.5: Blockchain with virtual RDF data and metadata

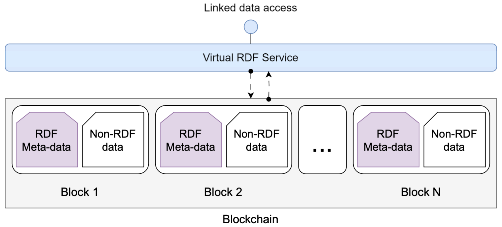

Figure 4.6: Blockchain with RDF metadata and virtual RDF

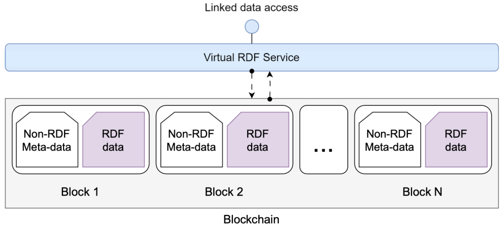

Figure 4.7: Blockchain with RDF data and virtual RDF

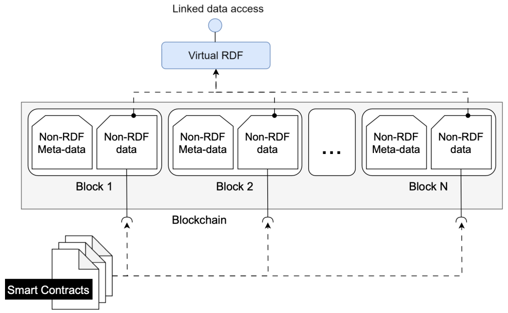

Figure 4.8: Blockchain with data into the block using smart contracts

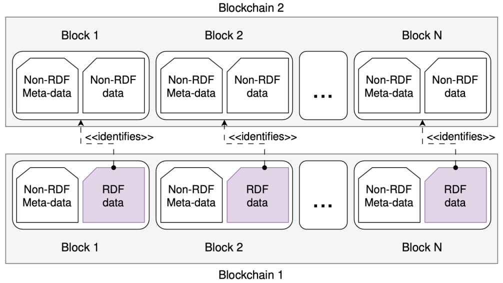

Figure 4.9: Blockchain referencing another blockchain with semantic data

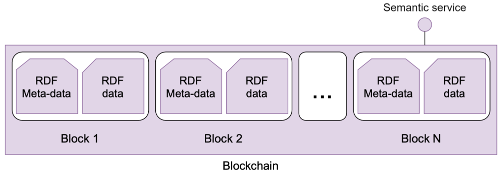

Figure 4.10: Blockchain implementation relying on semantic web technologies

Cano-Benito J.; Cimmino A.; García-Castro R. Towards Blockchain and Semantic Web. 2nd Workshop on Blockchain and Smart Contract Technologies (BSCT 2019). Workshops. LNBIP 373, pp. 220-231. Springer. Sevilla, España. 26-28 June 2019

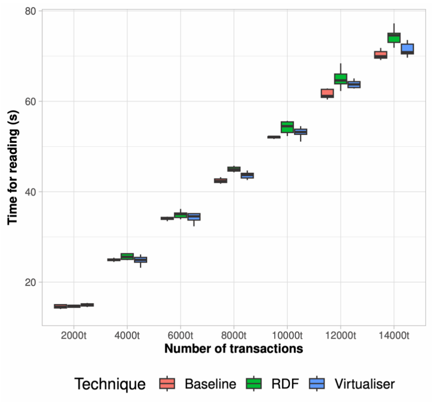

Figure 5.4: Time for reading transactions

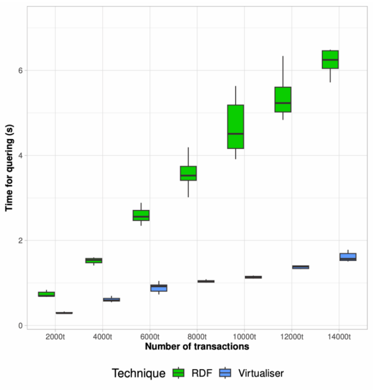

Figure 5.5: Time for querying data in the blockchain

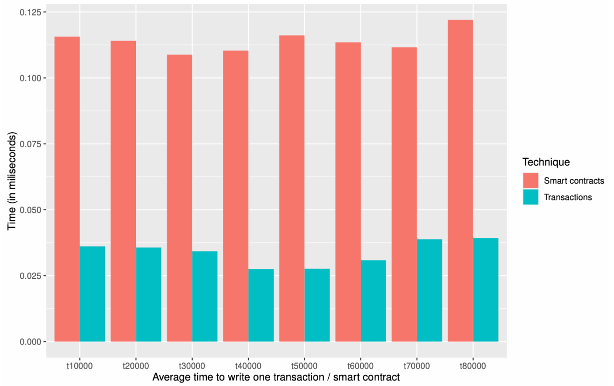

Figure 5.10: Average time to write one transaction / smart contract

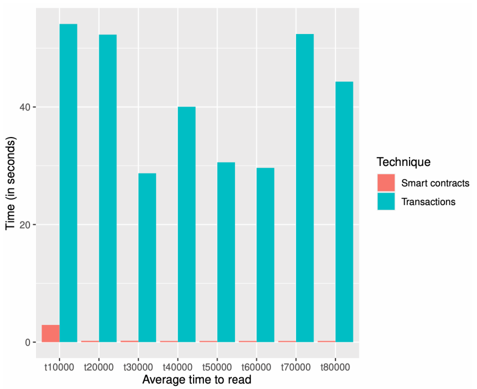

Figure 5.11: Average time to read one transaction / smart contract

- Should I store RDF directly in the blockchain? In which serialization?
  - Storing Turtle in a blockchain has some drawbacks:
    - Writing has an elevated cost in terms of gas
    - Reading takes more time than reading in JSON
  - Virtualising JSON stored in a blockchain is more efficient than storing Turtle
- Would it be better to use smart contracts to store RDF into the blockchain?
  - Storing data via smart contracts may involve a higher cost (gas) compared to transactions
    - It depends on the specific code used and the nature of the data being stored, e.g., arrays
  - Storing JSON-LD via transactions or via smart contracts?
    - Writing JSON-LD via transactions is more efficient
    - Reading from smart contracts is substantially faster

Cano-Benito J.; Cimmino A.; García-Castro R. Benchmarking the efficiency of RDF-based access for blockchain environments. 32nd International Conference on Software Engineering and Knowledge Engineering (SEKE 2020). Pp. 554-559. Pittsburgh, USA. 9-11 July 2020

Cano-Benito J. Convergence and representation of blockchain and smart contracts using the semantic web. PhD Thesis. TBD

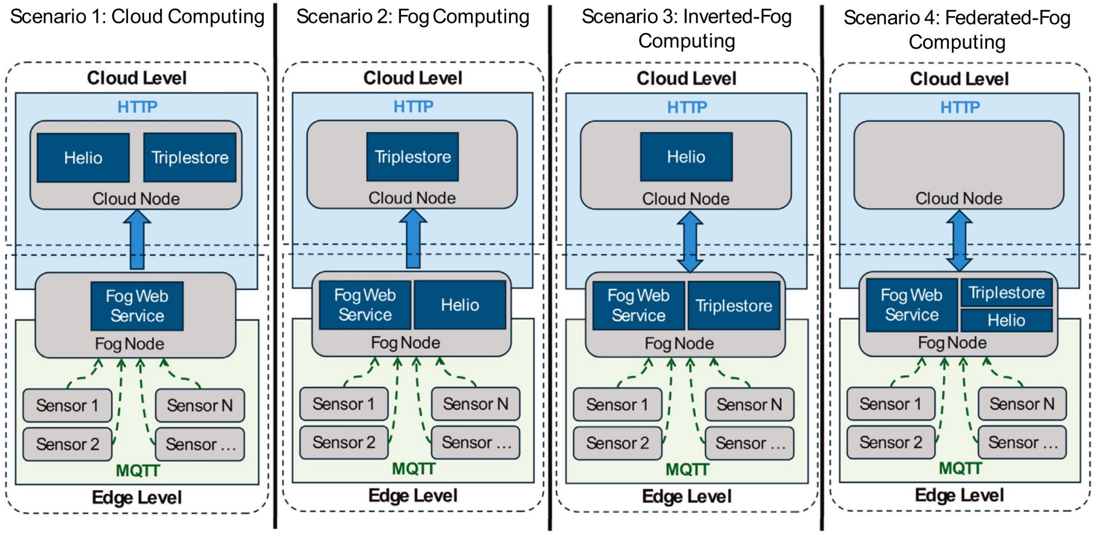

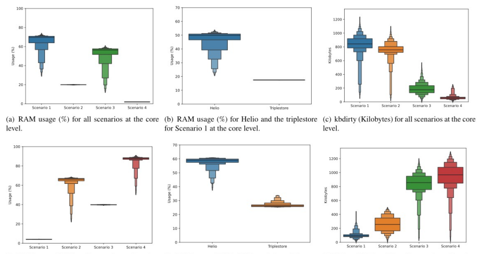

(a) RAM usage (%) for all scenarios at the core level. (b) RAM usage (%) for Helio and the triplestore for Scenario 1 at the core level. (c) kbdirty (Kilobytes) for all scenarios at the core level.

(d) RAM usage (%) for all scenarios at the edge level. (e) RAM usage (%) for Helio and the triplestore for Scenario 4 at the edge level. (f) kbdirty (Kilobytes) for all scenarios at the edge level.

Fig. 7. RAM parameters in the stress test at both the core and edge levels.

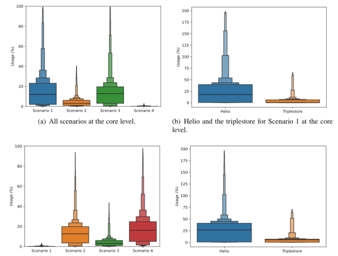

(a) All scenarios at the core level. (b) Helio and the triplestore for Scenario 1 at the core level.

(c) All scenarios at the edge level. (d) Helio and the triplestore for Scenario 4 at the edge level.

Fig. 6. CPU Usage (%) parameter in the stress test at both the core and edge levels.

Table 2
Definition and representation of performance parameters at the core level.

Parameter	Scenario 1	Scenario 2	Scenario 3	Scenario 4
CPU	Stable	not critical	Stable	Not critical
RAM	Stable	Not critical	Stable	Non-critical
kbdirty	Stable	Stable	Non-critical	Non-critical
TPS	Stable	Stable	Non-critical	Non-critical
KB_read	Non-critical	Non-critical	Non-critical	Non-critical
KB_write	Stable	Stable	Non-critical	Non-critical
rxpck	Stable	Non-critical	Stable	Non-critical
txpck	Stable	Non-critical	Stable	Non-critical
Translation	Stable	Stable	Stable	Stable
Transaction	Stable	Stable	Stable	Stable

Table 3
Definition and representation of performance parameters at the edge level.

Parameter	Scenario 1	Scenario 2	Scenario 3	Scenario 4
CPU	Non-critical	Critical	Non-critical	Critical
RAM	Non-critical	Critical	Stable	Critical
kbdirty	Non-critical	Non-critical	Stable	Stable
TPS	Non-critical	Non-critical	Stable	Stable
KB_read	Non-critical	Non-critical	Non-critical	Non-critical
KB_write	Non-critical	Non-critical	Stable	Stable
rxpck	Stable	Stable	Stable	Stable
txpck	Stable	Stable	Stable	Stable

- It is key to perform strategic resource allocation and performance optimisation in architectures
  - No critical issues at the cloud level; the edge level requires meticulous software assessment
- Where to focus on resource-constrained systems?
  - CPU and RAM consumption (and only for data translation; no problem with storage)
  - Network, latency and hard disk load demonstrate stability and non-critical behaviour
- Federation is promising:
  - Optimises resource utilisation and enhances efficiency
  - Enables federated machine learning and privacy

Huaranga-Junco E.; González-gerpe S.; Castillo-Cara M.; Cimmino A.; García-Castro R. From Cloud and Fog Computing to Federated-Fog Computing: A Comparative Analysis of Computational Resources in Real-Time IoT Applications based on Semantic Interoperability. Future Generation Computing Systems. Pre-print.

Conclusions

Conclusions regarding semantics

- Everything is about reaching agreements…
  - Think of outsiders
- …and sharing!
  - More than just the implementation
- Semantic interoperability goes beyond just using one ontology
- Your use of semantics (e.g., expressivity) will change over time
  - How will others know?

Conclusions regarding interoperability

- Things frequently overlooked
  - Heterogeneity is (and will be) a reality
  - Interoperability is social and technological
- No golden bullet for semantic interoperability
  - Engineering semantic interoperability in an evolving context
- Where to put the intelligence?
  - Need experimental body of knowledge
- Full interoperability is still a challenge
  - Automation
  - Certification
  - Governance
  - …

Questions

rgarcia@fi.upm.es

LDAC 2024

Ontology Engineering Group

Lessons learnt from researching on semantic interoperability

Raúl García-Castro

Acknowledgements: Funded by the Madrid Government (Comunidad de Madrid-Spain) under the Multiannual Agreement with the Universidad Politécnica de Madrid in the Excellence Programme for University Teaching Staff, in the context of the V PRICIT (Regional Programme of Research and Technological Innovation).

Comunidad de Madrid
Dirección General de Investigación e Innovación Tecnológica
CONSEJERÍA DE CIENCIA, UNIVERSIDADES E INNOVACIÓN

rgarcia@fi.upm.es

13/06/2024

Bochum