

Standard International Trade Classification

From Spreadsheet to OWL-2 Ontology

Norbert Luttenberger · Jesper Zedlitz

Received: 19 December 2016 / Accepted: 10 April 2017 / Published online: 1 September 2017
© Springer Fachmedien Wiesbaden GmbH 2017

Abstract Trade classifications are a necessary prerequisite for the compilation of trade statistics, and they should – beyond that – be regarded as a valuable base for the definition of shared controlled vocabularies for linked business data that deal with import, export etc. The Standard International Trade Classification (SITC) provided by the UN Statistics Division is a widely used classification mostly applied for scientific and analytical purposes. SITC – as most other trade classifications – is available today only in text or spreadsheet formats. These formats reveal the inner hierarchical structure of SITC to the human reader, because SITC trade codes are built according to the decimal classification scheme, but unfortunately, SITC’s inner structure is opaque to computer applications in text and spreadsheet formats. The paper discusses an approach to set up an OWL-2 ontology for SITC that states subsumption relations between classes of goods. This kind of semantic underpinning of SITC is suited to ease both checking and extending SITC and to derive from it a shared controlled vocabulary for business linked data. Some problems of today’s SITC (among them missing inner nodes of the trade code hierarchy) are carefully discussed, and the paper motivates several decisions that were taken for ontology design. Finally, the study introduces the semantic reasoner as a tool for the (at least partial)

automatic derivation of structural information for SITC from the trade code building rule. The paper reports on reasoner runtimes observed for different versions of the SITC ontology and for different versions of the Pellet reasoner.

Keywords Trade classification · Semantic reasoning

1 Introduction

Trade classifications are a necessary prerequisite for the preparation of trade statistics that are used to describe – for administrative and/or scientific purposes – domestic and international flows of goods. Without challenging this intended use, we discuss trade classifications in this paper with a different motivation: We argue that trade classifications – beyond their obvious purpose – can also be considered as valuable sources for the definition of shared controlled vocabularies. Shared controlled vocabularies are at the foundation of Linked (Open) Data collections. With this perspective in mind, we argue that the terminology work that is required for the provision of meaningful linked business data can profit from existing vocabularies, among them those that give trade statistics their shape.

Unfortunately, most trade classifications today are available in text or table (spreadsheet) formats. These formats address the human reader; they are not well suited to reveal the inner, mostly hierarchical structure of trade classifications to computer processing and examination. Experience has shown that with this kind of formats structural problems may arise on several occasions, for instance when adjusting a trade classification to new demands. To avoid these problems, we therefore prefer a logics-based format for the development of shared

Accepted after 2 revisions by Prof. Dr. Abramowicz.

Prof. Dr. N. Luttenberger (✉) · Dr. J. Zedlitz
Research Group for Communication Systems, Department of
Computer Science, Christian-Albrechts-University in Kiel, Kiel,
Germany
e-mail: n.luttenberger@email.uni-kiel.de

Dr. J. Zedlitz
e-mail: j.zedlitz@email.uni-kiel.de

controlled vocabularies, i.e. a format that allows us to ground “vocabulary control” on logical reasoning (and related tools).

A second argument is in favor of our proposition: Today, we see numerous trade classifications (see below), some of them covering (slightly) different purposes, some of them covering different historical epochs. It would be an endeavoring task to find out which parts of such classifications are equivalent and which parts are not. Though this task is beyond the scope of this paper, we feel free to point out that in text based formats, checking for equivalence has only string comparison as basic operation, while in logics-based formats, we can additionally take structural properties into account.

Against this backdrop, we see our effort as an “example-driven” feasibility study for the derivation of shared controlled vocabularies from trade classifications. The paper focusses on the development of an OWL-2 ontology (term defined below) for the fourth revision of the *Standard International Trade Classification* (SITC-4) that was published in 2006 by the United Nations Statistics Division (UNSD) (United Nations Statistics Division 2010). SITC-4 – currently published in text and spreadsheet formats – plays an important role mostly for analytical purposes in the economics area. The UNSD states: “Many countries and national and international organizations continue to use SITC for various purposes, such as for the study of long-term trends in international merchandise trade and aggregation of traded commodities into classes more suitable for economic analysis.” (United Nations Statistics Division 2015). Similar to the SITC-4 with respect to its hierarchical structure is the *Combined Nomenclature* (CN) (European Commission 2015); it is used in the European Union mostly for administrative purposes. The CN builds upon and extends the so-called *Harmonized System*, which has been developed by the World Customs Organization (2015). In this paper, we also take into account the correspondence table that defines mappings between SITC-4 and CN codes (Eurostat 2015a, b) (called “correspondence table” for short in the remainder of this paper). The manifold applications and the worldwide usage of both SITC-4 and CN justify the effort to improve their structure representation.

There exist numerous definitions for the term “ontology”. Studer et al. give a widely accepted definition: “An ontology is a formal, explicit specification of a shared conceptualization.” (Cited in Staab and Studer 2009). Guarino, Oberle, and Staab explain: “A conceptualization is an abstract, simplified view of the world that we wish to represent for some purpose.” (Staab and Studer 2009). ISO 1087-1:2000 defines “concept” as a “unit of knowledge created by a unique combination of characteristics.” (International Organization for Standardization 2000). In

general, ontologies thus can be seen as collections of knowledge for some specific domain of discourse.

More specifically, formal ontologies (e.g., OWL-2 ontologies) are ontologies that are built upon description logics (Krötzsch et al. 2013). Here, “concepts” are mapped to classes, and “characteristics” are mapped to properties. A formal ontology (“ontology” for short in the remainder of this paper) deals with *classes*, class subsumption, and class relationships. OWL-2 offers a wealth of class expressions that impose necessary and/or sufficient conditions for class membership. Besides classes, OWL-2 allows for the description of *named individuals*. Named individuals carry data and/or object properties and are assigned to classes explicitly or by a reasoning process. Ontology statements are called *axioms*. Within an ontology, we can distinguish between the TBox and the ABox. TBox axioms describe concepts (e.g., extensional or intensional definition of classes, class/subclass relations), while ABox axioms capture knowledge about named individuals (e.g., assign individuals to classes) (Krötzsch et al. 2013). Most important when discussing ontologies is the reasoning support that comes with ontologies. Reasoning – among other things – detects “hidden knowledge” in ontologies, for instance class subsumption or class membership (when these facts have not been written down by the ontology engineer) and checks for ontology consistency.

OWL-2 ontologies can be used as semantically rich representations of class hierarchies, and therefore they lend themselves very well to being applied as format for trade classifications. The reasoning support coming with ontologies helps to detect inconsistencies and structure problems in trade classifications. To a certain extent, reasoning can also be seen as a tool for ontology generation, or in our view: a tool for generating structurally correct trade classifications from partial descriptions.

Though at first sight the approach chosen in this paper seems to be straight forward, it turned out during our project that two severe problems had to be solved: First, in both SITC-4 and the correspondence table, several structure problems had to be patched before a formal ontology could be developed. Second, reasoner runtimes for the initial version of our ontology were very long, i.e. several hours. We re-coded parts of the open-source Pellet reasoner to enhance parallel program execution, and we re-engineered our ontology. For a detailed discussion of the influence of different choices we made, the reader is referred to Sect. 6 of this paper.

Our paper is organized as follows: In the following section, we discuss approaches similar to ours but applied to other trade classifications like, for instance, the United Nations Standard Products and Services Code (UNSPSC), or the eCI@ss cross-industry classification of products and services. In Sect. 3 we give a short introduction to SITC-4

and its problems. Section 4 explains in depth our design decisions for transforming SITC-4 into an ontology. In Sect. 5, we briefly present the transformation workflow. A special focus is put on reasoning in the following section. The paper closes with a short summary and an outlook on further work.

2 Related Work

Tolksdorf et al. (2003) give a general overview over semantic technologies for the business area. The authors identify “three important building blocks” for e-commerce scenarios, namely “the use of URIs as a global identification mechanism for products and traders, the RDF data model, and the Web Ontology Language (OWL) for the definition of common terms and concepts”.

In a similar intention, Ding et al. (2004) discuss the “The Role of Ontologies in eCommerce”. From their perspective, ontologies may be a prerequisite for a common understanding of product catalogs in B2B transactions. In this context, they also discuss the UNSPSC trade classification (see below).

Unfortunately, both papers do not cover in-depth technical details. In the following, we therefore concentrate on papers that deal with trade classifications and with technical approaches to make their inherent structure explicit. The material is organized in a sequence that shows increasing formal effort for superimposing structural views on existing trade classifications.

We start with a historic trade classification that does not even use the decimal classification scheme for trade code building. In 2014, the Deutsche Zentralbibliothek für Wirtschaftswissenschaften – Leibniz-Informationzentrum Wirtschaft (ZBW) published the “Statistik des Deutschen Reichs 1873–1883” (*Reichsstatistik*) on Open Access (Deutsche Zentralbibliothek für Wirtschaftswissenschaften 2015a). The statistical data included in the 10 volumes of the *Reichsstatistik* refer to a historic trade classification that comprises approx. 450 terms for trade goods. Terms are ordered in a very elementary 2-layer hierarchy, but despite of this hierarchical ordering the *Reichsstatistik* does not attempt to assign hierarchical codes to the listed goods. ZBW has published the *Reichsstatistik* in HTML, Excel, and other formats. A structured view on the *Reichsstatistik* trade classification has not been developed.

As part of its *Standard Thesaurus Wirtschaft* (STW – Thesaurus for Economics), ZBW has published a sub-thesaurus for “commodities” that includes a number of entries for products of different kinds. The ZBW has published the STW in different formats, among them RDF and Turtle. The STW uses SKOS (Simple Knowledge Organization System) (W3C 2012) to structure its thesaurus. SKOS

provides a vocabulary that puts terms of domains of interest into relation to each other; such relations are, for instance, “broader term”, “narrower term”, and “related term”. In doing so, SKOS follows ISO 25964, the international standard for information retrieval thesauri. ISO 25964 defines a thesaurus as a “controlled and structured vocabulary in which concepts are represented by terms, organized in a way that relationships between concepts are made explicit, and preferred terms are accompanied by lead-in entries for synonyms or quasi-synonyms” (International Organization for Standardization 2015). Though using SKOS adds structural elements to an otherwise flat catalogue of words and thus is adequate for setting up a thesaurus, we do not intend to follow this approach for two reasons:

- We think that a trade classification does not deal with terms, but that it deals with the classification of goods themselves. Such a classification establishes real subsumption relations between classes of goods (rdfs:subClassOf). “Terms” play only a minor role in a classification, they are used mainly for creating annotations that can be accessed by the human reader.
- SKOS aims at poly-hierarchies, i.e., a term may be, for instance, a “narrower term” to more than one “broader term”. A trade classification in contrast is a strict mono-hierarchy, where each trade code belongs to only one class in the next higher layer. Enforcing a mono-hierarchy in SKOS requires additional complexity that makes the ontology less comprehensible.

Yet another direction is chosen by eCI@ss, a “cross-industry product data standard for classification and description of products and services” with “applications such as procurement, controlling and distribution” (eCI@ss e.V. 2015). It has “40,800 product classes and 16,800 properties” covering “the majority of traded goods and services”. In eCI@ss, products and services are described in the OntoML language. OntoML stands for “Product Ontology Mark-up Language” and is standardized in ISO 13584-32. OntoML is an XML schema designed for use by applications that need to exchange and process ISO 13584 PLIB (Parts LIBrary) compliant domain ontologies, possibly together with their related instances, in various Web-oriented environments. This schema – as any other XML schema – uses the element nesting capabilities provided by XML to express relations between XML elements. This is an obviously valid notation for hierarchical structures. But OntoML instance documents can only be checked syntactically by a validating parser, and not semantically, because XML Schema is not based on description logics as, for instance, OWL-2, where one can apply a reasoner to check and even establish a correct subsumption of classes.

Hepp and de Bruijn (2007) describe their attempt to develop a generic methodology for deriving OWL and RDF-S ontologies from “hierarchical classifications, thesauri, and inconsistent taxonomies”. They argue that “informal hierarchical categorizations” may apply inter-class relations that are not subsumptions, and therefore they introduce a so-called “context” that installs the set of relations needed in the respective situation. They illustrate their approach by a fictitious hierarchy, where “ice cubes” are a subcategory (but not a subclass!) of “beverages”, and argue that this might be the perspective of a “purchasing manager”. They exemplify their very generic methodology by the construction of OWL ontologies for both eCI@ss (see above) and the United Nations Standard Products and Services Code (UNSPSC 2014). In contrast to Hepp and de Bruijn (2007), we focus on classifications with subsumption relations only, and thus have the benefit that the reasoner can fill in missing subsumption relations automatically, i.e., we do not need what Hepp and de Bruijn call a context.

Stolz et al. (2014) present a tool called PCS2OWL that according to the authors allows the user to transform trade classifications from various formats into OWL ontologies. Unfortunately, their approach seems to be overly complex: The mentioned tool creates two OWL classes from each category in the source product classification. “The first is a broader taxonomic class that represents the category from the [product classification] in the target ontology. The second is a context-specific class, in our case in the domain of products and services.” This obviously doubles the number of concepts in the resulting ontology. Another severe concern is that in the approach of Stolz et al. the resulting ontology is not checked by a reasoner, but only by a number of SPARQL queries, comparing category and class counts. Thus, structural issues are beyond the scope of inspection.

Finally, Caracciolo et al. (2012) show the development of an OWL ontology for the fisheries domain. The ontology is built upon FAO’s International Standard Statistical Classification of Fishery Commodities (ISSCFC) (2002), which is an expansion of the SITC, and it is linked with the Harmonized System. Nevertheless, their ontology seems to be an ontology “from scratch”, it is not just a “snippet”

from SITC-4. It nevertheless is linked to SITC-4 by providing a per-item property that holds the related SITC-4 trade code.

3 Short introduction to SITC-4 and its Problems

SITC-4 can be modeled as a monohierarchical classification with 5 hierarchy levels, called *tiers*. Below the unnamed root we find – in downward order – *sections* (tier 1), *divisions* (tier 2), *groups* (tier 3), and *subgroups* (tier 4). Finally, tier 5 contains so-called *basic headings*. SITC-4 has 10 sections, 66 divisions, 262 groups, 1023 subgroups, and 2652 basic headings. In the following, we are going to explain why it is advisable to add further basic headings to SITC-4.

SITC-4 trade codes are formed according to the basic rules of a decimal classification. [When referring to “basic” we exclude auxiliary signs and auxiliary numbers as used, for instance, in the so-called Universal Decimal Classification (UDC Consortium 2015)]. More formally, a trade code TC is defined as follows:

$$0 \leq c_i \leq 9, \quad c_i \in \mathbb{N}$$

$$1 \leq i \leq 5, \quad i \in \mathbb{N}$$

$$TC ::= c_1|c_1c_2|c_1c_2c_3|c_1c_2c_3c_4|c_1c_2c_3c_4c_5.$$

Sections have a 1-digit code, divisions have a 2-digit code, and so on. The first digit of a trade code identifies the related section, the first two digits identify the related division, and so on. In the following, we call the single digits c_i of a trade code its tier-1 code, its tier-2 code, and so on. A sample “path” through the SITC-4 hierarchy is given in Table 1.

A deeper analysis of the SITC-4 shows that the classification has 318 subgroups that have no subordinate basic headings. In the following, we call these subgroups “left-alone subgroups”.

In order to provide basic headings to the left-alone subgroups, we consulted the correspondence table. Indeed, the correspondence table contains 317 basic headings that are suited to fill the gaps below the left-alone subgroups. Nevertheless, one remaining subgroup keeps its status as left-alone subgroup. Figure 1 shows the result of the

Table 1 Sample SITC-4 trade codes

	c_1	c_2	c_3	c_4	c_5	
Section	0					Food and live animals
Division	0	5				Vegetables and fruit
Group	0	5	7			Fruit and nuts (not including oil nuts), fresh or dried
Subgroup	0	5	7	1		Oranges, mandarins, clementines and similar citrus hybrids, ...
Basic Headings	0	5	7	1	1	Oranges, fresh/dried
	0	5	7	1	2	Mandarins (including tangerines and satsumas); clementines, ...

Fig. 1 Combining SITC-4 and correspondence table trade codes

SITC-4	basic headings	2652			
	left-alones	317			1
Corr. table	basic headings	2979			
	Orphans	17			
combined	basic headings	2986			
	not mapped				7

Table 2 Orphan trade codes and related gaps in the SITC-4 hierarchy

codes	60				70		80	94		99
missing in SITC-4	600		660		700		800	941	972	998
	6000	6580	6600	6950	7000	7910	8000	9410	9720	9988
orphan basic headings from corr. table	60001	65800	66000	69500	70001	79100	80001	94100	97200	99888
	60002				70002		80002			
	60003				70003					
	60004									
	60005									

combination of SITC-4 and the correspondence table (bar length not proportional).

When taking over all non-duplicate trade codes from the correspondence table into a combined classification, we are confronted with a new phenomenon: The correspondence table delivers 17 basic headings that are not related to any subgroup in the “original” SITC-4. For some of these basic headings not only the subgroup layer is missing, but also “ancestor” layers further up in the hierarchy. We call these basic headings “orphans”. Table 2 shows the orphans together with the related division, group, and subgroup codes that are needed to fully populate the hierarchy. The German version of the SITC-4, provided by the German Statistisches Bundesamt (Federal Statistical Office), mentions that some of these orphans have been introduced as national specifics (Statistisches Bundesamt 2006).

The analysis given above clearly indicates that it is hardly recommendable to convert SITC-4 on an as-is basis to an OWL-2 ontology. In following chapter, we show how we dealt with the identified problems.

4 Approach

In the following subchapters, we explain our design guides for the development of an OWL-2 ontology for the SITC-4. These design guides seem very specific, but could serve “best practice” rules for solving similar problems.

4.1 Orphan Basic Headings and Left-alone Subgroups

As mentioned above, inserting trade codes of the correspondence table into SITC-4 produced orphan basic headings. With our goal in mind to make the classification structure fairly regular, we had two choices: Either prune these orphan codes, or generate artificial “ancestors”, i.e. subgroups, groups, and even divisions as detailed in Table 2. Even it is questionable to add this kind of pseudo-information (i.e. SITC-4 codes that do not stand for explicitly defined product groups) to SITC-4, we decided to generate ancestors for orphans, i.e. codes that play the role of stopgaps. If we had decided otherwise, we had lost several trade codes. Generating these artificial ancestors is easy as we can deduce their codes from the codes of the orphan basic headings.

Even after inserting trade codes of the correspondence table into SITC-4, one left-alone subgroup remained in the SITC-4. With our goal in mind to make the classification structure fairly regular, we could have considered to implement a “dummy” basic heading such that it can be subsumed under this left-alone subgroup. In fact, we decided to leave this left-alone subgroup as it is, which means that it adopts the role of a product code, similar to a basic heading. This decision was motivated by our unwillingness to add more information items than absolutely necessary to the “original” SITC-4.

4.2 SITC-4 Basic Headings

At first sight, it might seem rather “natural” to map all members of the five tiers of SITC-4 to OWL-2 classes and to not consider any class instances (in OWL-2 terminology: named individuals) in a trade classification. This procedure would have left the usage of the OWL-2 language construct `owl:NamedIndividual` to denote concrete instances of some product, i.e. “this box of oranges”. However, we followed another approach: We mapped the members of SITC-4 tiers 1–4 to OWL-2 classes (`rdf:type owl:Class`), and we mapped the SITC-4 basic headings to OWL-2 named individuals (`rdf:type owl:NamedIndividual`). The reason for this design decision can be summarized as follows:

We consider SITC-4 as an instrument for the preparation of trade statistics. We assume that trade statistics are built from datasets combining e.g. product kind, value, time span, and additional information. In RDF, such a dataset (in a sample namespace abbreviated “stat”) could be written (in Turtle syntax) as a set of triples for instance like: `stat:exportItem stat:ofKind sitc4:05711; stat:hasValue “1.5 M USD”; stat:inYear “2010”`. In this sample dataset, the trade code 05711 (see Table 1) is used as value of an object property. For object property values, only individuals can be used in an ontology. If we had decided to map SITC-4 basic headings to OWL-2 classes, the designer of an ontology for trade statistics would have been in danger to write down the sample `stat:exportItem` dataset as follows: `stat:exportItem rdf:type sitc:05711 ...` A reasoner would have concluded from this modeling that `stat:exportItem` datasets were of type “Oranges, fresh/dried” which is obviously nonsense.

A second argument is valid, too. Modeling the SITC-4 basic headings as OWL-2 individuals allows us to assign further (data/object) properties to these individuals. In the next chapter, we show a sample property assignment by linking additional information from the correspondence table (CN trade code and validity period) to SITC-4 basic headings.

Find a similar discussion on the topic “class vs. individual” in Noy and McGuinness (2001).

4.3 CN Trade Codes and Their Validity Periods

The correspondence table maps almost all SITC-4 basic headings to related CN trade codes, and additionally gives a validity period for these mappings. As can be seen from the correspondence table, SITC-4 trade codes have been mapped to different CN trade codes over time.

In order to enrich our ontology with CN trade codes and related validity periods, we need some kind of (ad-hoc) reification for statements that link a CN trade code to a SITC-4 trade code. The W3C Working Group Note on

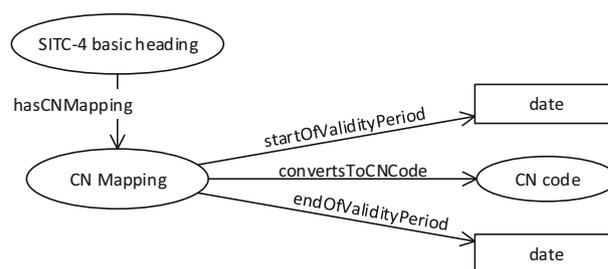


Fig. 2 CN Mapping: attaching a CN code and its validity period to a SITC-4 basic heading via a UUID-identified node

“Defining N-ary Relations on the Semantic Web” (Noy and Rector 2006) defines in its Use Case 1 a pattern that can easily be applied to our problem as well. [For a more in-depth discussion on “Time-Dependent Factual Knowledge”, see Krieger and Declerck (2015)]. We constructed an ad-hoc reification for the SITC-4-to-CN mapping (cf. Fig. 2).

In a “pure” RDF environment, one could link reifying statements to a node via an RDF blank node. To enable a distinct access to any property combination, we decided to not apply the RDF blank node construct. Instead, we combine the properties `:convertsToCNCode`, `:startOfValidityPeriod`, `:endOfValidityPeriod`, and their values by named individuals that are identified by Universally Unique Identifiers (UUID) according to ITU-T Rec. X.667, which are basically pseudo-random numbers. We call these individuals “CN Mappings” in the remainder of this paper. CN Mappings are related to SITC-4 basic headings by the object property `:hasCNMapping`. CN Mappings are collected in the class `:CNMapping` (not shown in Fig. 2).

4.4 Inferred Class Subsumption

A reasoner may – at least partially – generate subsumption relations, assuming that we define suited class axioms. Thus, we have the choice to either include explicit class subsumption axioms into our ontology (using `rdfs:subClassOf` relations), or to construct class expressions such that class subsumption can be inferred by the reasoner automatically. To ease the insertion of new sections, divisions, groups and subgroups into SITC-4 or the deletion of such from SITC-4 or the modification of the SITC-4 structure, we think that SITC-4 tiers should be as independent from each other as possible. This prohibits the extensive use of `rdfs:subClassOf` relations. We therefore state our goal as follows: Class subsumption between SITC-4 sections, divisions, groups and subgroups is to be inferred automatically by the reasoner; the ontology should be constructed such that it is not required to state any `rdfs:subClassOf` relation between members of different

```

1  :hasTier1Code rdf:type owl:DatatypeProperty , owl:FunctionalProperty .
2  :hasTier2Code rdf:type owl:DatatypeProperty , owl:FunctionalProperty .
3  :hasTier3Code rdf:type owl:DatatypeProperty , owl:FunctionalProperty .
...
4  :S246      rdf:type owl:Class ;
5            owl:equivalentClass [ rdf:type owl:Class ;
6                                  owl:intersectionOf (
7                                    [rdf:type owl:Restriction ; owl:onProperty :hasTier1Code ;
8                                      owl:hasValue "2" ]
9                                    [rdf:type owl:Restriction ; owl:onProperty :hasTier2Code ;
10                                   owl:hasValue "4" ]
11                                   [rdf:type owl:Restriction ; owl:onProperty :hasTier3Code ;
12                                   owl:hasValue "6" ]
13                                 )
14          ]
15 .

```

Code snippet 1 Necessary and sufficient conditions for class membership

tiers. From a more general perspective, we could call this “inferring knowledge from data”.

We proceed as follows (see example in code snippet 1, Turtle syntax): We define functional datatype properties :hasTier1Code, ..., :hasTier5Code (see lines 1–3). Next, we define necessary and sufficient conditions for class membership (owl:equivalentClass) by restricting the values of these properties (owl:hasValue) and by building an intersection of these properties, lines 4–15 of code snippet 1.

Now, we can leave it to the reasoner to find the correct positions for all classes in the class hierarchy: Classes with less conditions set are superclasses of classes with more conditions set, assuming that the property values match.

4.5 Class Disjointness

SITC-4 is constructed such that SITC-4 sections do not share any SITC-4 divisions, SITC-4 divisions do not share any SITC-4 groups, and SITC-4 groups do not share any SITC-4 subgroups. In other words, the SITC-4 trade code hierarchy is not a poly-hierarchy. To represent this SITC-4 construction principle in an OWL-2 ontology we have to take care for “tier-wise” disjointness of all OWL-2 classes representing SITC-4 sections, divisions, groups, and subgroups.

We could express this by the OWL-2 language construct owl:AllDisjointClasses. But we decided to leave the detection of class disjointness to the reasoner:

- We made the :hasTier1Code, ..., :hasTier5Code properties functional. This means that to each individual

basic heading at most one distinct tier 1, ..., 5 code can be assigned.

- We constructed the classes related to SITC-4 trade codes by necessary and sufficient conditions on the :hasTier1Code, ..., :hasTier5Code properties.
- We assigned different sets of tier 1, ..., 5 codes to the SITC-4 trade codes.

This construction enables the reasoner to detect class disjointness automatically.

4.6 Assignment of Individuals to Classes

A last question remains: Can we leave it to the reasoner to assign SITC-4 basic headings to SITC-4 subgroups? Or in ontological terminology: Can we leave it to the reasoner to assign individuals to classes? Obviously, the rdf:type language construct, which is part of OWL-2, would allow us to explicitly assign individuals representing SITC-4 basic headings to SITC-4 classes. But following the idea that a reasoner should be used not only as an instrument for checking a classification, but also as an instrument for partially generating a classification, we would prefer to leave the assignment of individuals to classes to the reasoner. Automatic assignment exploits the :hasTier1Code, ..., :hasTier5Code properties and the value sets assigned to them, because they “work” both as parts of class expressions and as properties for individuals. We re-discuss this problem in more depth in chapter 6 of this paper.

5 From Excel Sheet to OWL-2 Ontology

Sources for the transformation of SITC-4 to an OWL-2 ontology are both the SITC-4 Excel sheet and the correspondence table Excel sheet. A large number of tools is available to transform from Excel format to RDF format; a comprehensive list of tools for this purpose is given in W3C (2015). We decided against using any of these tools for the following reasons:

- In our case, we do not only have to solve a conversion problem, but also a combination problem.
- We do not simply transform from Excel to RDF format, but from Excel format to OWL “format”.
- Some of the listed tools assume the existence of a database scheme that may serve as a conversion aid. For SITC-4 and the correspondence table such a schema is not available.

We decided to develop a dedicated tool chain that is mainly based on XSL(T). Our experience shows that this tool chain can easily be modified to carry out the intended format conversion also for other trade classifications. The tool chain is explained in the following.

After some preprocessing – splitting of SITC-4 trade codes into tier codes, converting validity dates into ISO 8601 format, converting table content from native Excel format to xml format, and other – the transformation is carried out in three steps (white boxes in Fig. 3).

1. Both Excel sheets are combined by an XSL(T) transformation (SAXON XSL(T) processor), i.e. CN Mappings are added to the ontology. Additionally, left-alone SITC-4 subgroups are equipped with basic headings from the correspondence table, and “ancestor” elements are generated for orphans from the correspondence table. The output is a file in XML format and denoted by “SITC-4+.xml” in Fig. 3.
2. The output of step 1 is transformed into an OWL-2 ontology by a second XSL(T) transformation. The resulting output file is denoted by “SITC-4+.owl” in Fig. 3. OWL is rendered in Turtle syntax.
3. Finally, the output of step 2 is processed by the Pellet reasoner. The reasoner infers additional information that is stored together with its input in a file denoted by “SITC-4++.owl” in Fig. 3.

The SITC-4+ ontology comprises 1387 classes and 21,218 individuals in total. 18,231 (!) of these individuals are CN Mappings. Obviously, most SITC-4 basic headings are affected by numerous changes of the mapping CN trade code.

The next section discusses the third step of Fig. 3 in more detail.

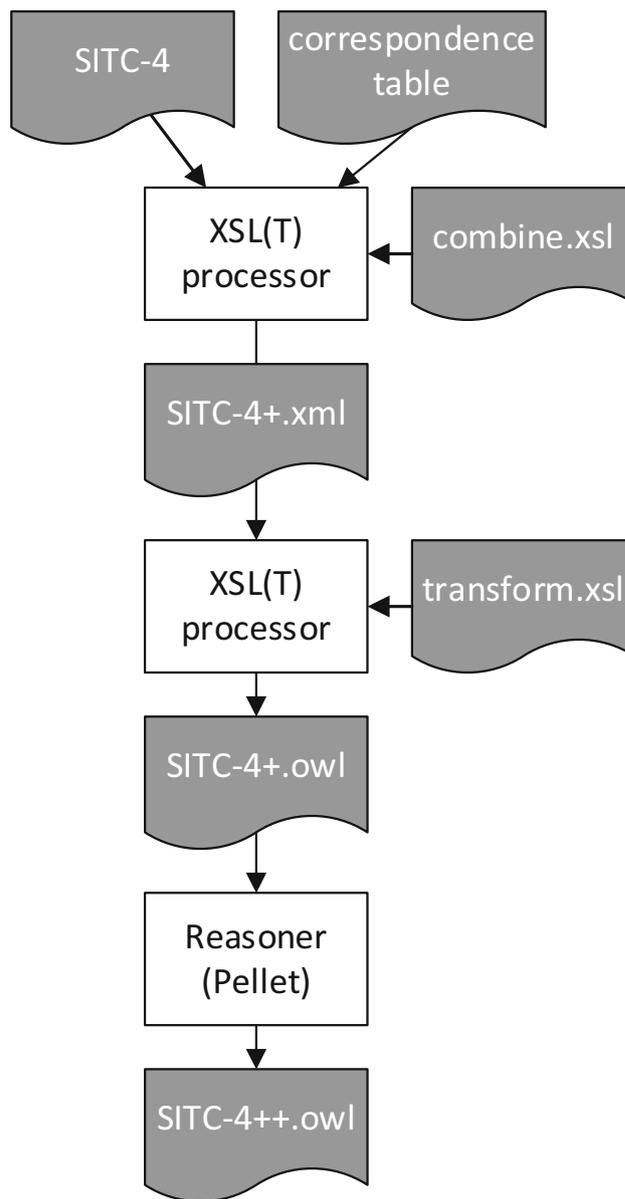


Fig. 3 Workflow

6 Reasoning

To set up the class hierarchy of the SITC-4+ ontology and to assign its individuals to classes, we applied the open-source version of the Pellet reasoner (version 2.3.1). The next chapter shows reasoning benefits in general. To reduce the initially observed reasoner runtimes, we derived from the “original” Pellet reasoner a variant that takes care of parallel assignment of individuals to classes (chapter 6.2). Additionally, we examined reasoner runtimes for three different ontology versions (chapter 6.3).

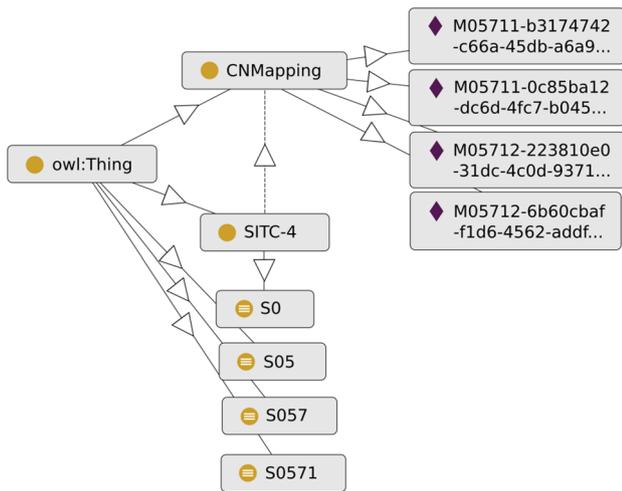


Fig. 4 SITC-4 ontology before reasoning

6.1 Reasoner as “Structure Generator”

Figures 4 and 5 below show two different graphical representations of a very small snippet from our ontology dealing with “Oranges, mandarins, clementines and similar citrus hybrids ...” (SITC-4 subgroup 0571). Both ontology snippets are visualized by the OntoGraf plugin of the ontology editor Protégé. Figure 4 shows the ontology

snippet before reasoning, Fig. 5 after reasoning. A box with a circle indicates a class, a box with a diamond indicates an individual. Arrows between classes show the existence of a subsumption relation or of a domain/range relation, arrows between classes and individuals show class membership, and finally arrows between individuals show object properties.

Two differences can easily be discovered: In Fig. 4, all classes except class S0 (standing for SITC-4 section 0) are direct subclasses of owl:Thing, while in Fig. 5, classes S0, ..., S0571 are in a subsumption relation, and only classes :CNMapping and :SITC-4 are direct subclasses of owl:Thing. The arrow between classes :SITC-4 and :CNMapping indicates that class :SITC-4 is the rdfs:domain of :CNMapping. Without reasoning, CN Mappings are assigned to a class (namely :CNMapping), and thus are rendered in the graphics, while after reasoning individuals representing SITC-4 basic headings, too, are assigned to classes, and thus are rendered. These individuals are members of classes SITC-4, S0, S05, S057, and S0571. Obviously reasoning adds a considerable amount of structuring information to the SITC-4+ontology.

Reasoning is a time-consuming task (see Table 3). Therefore, it is very worthwhile to carefully examine reasoner runtimes. In the following chapters, we describe

Fig. 5 SITC-4 ontology after reasoning

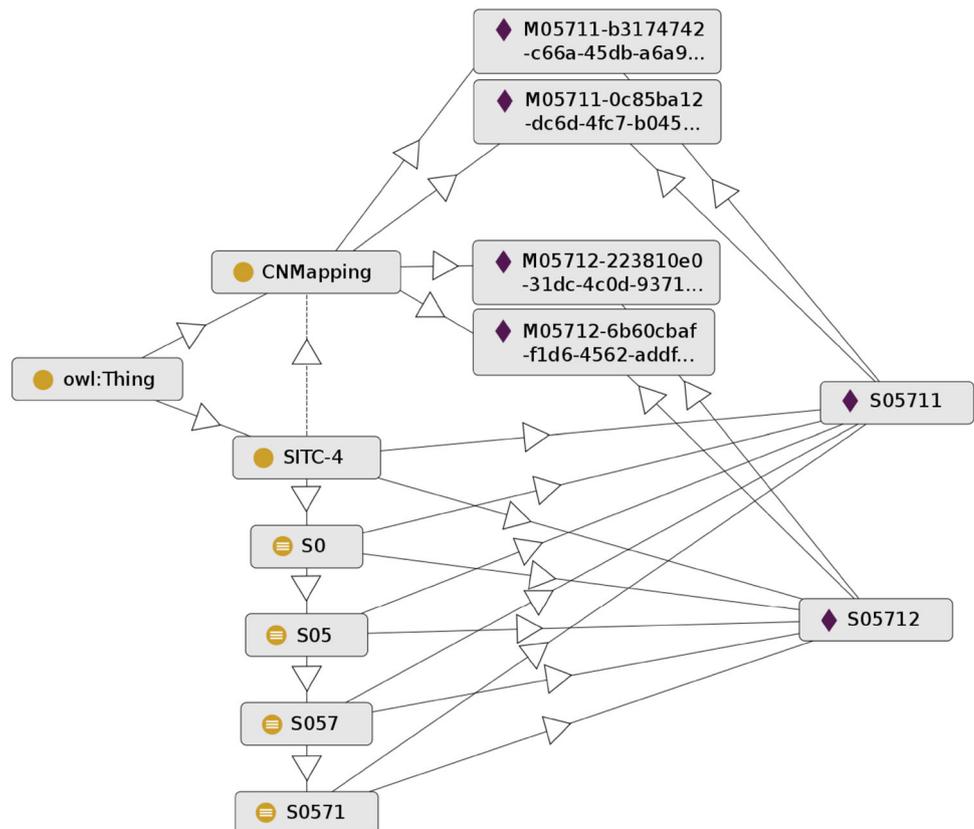


Table 3 Reasoner runtimes

	SITC-4+ _{impl}	SITC-4+ _{map}	SITC-4+ _{expl}
Sequential “Realizing”	3 h 10 min	3 h 15 min	6 min 3 s
Parallel “Realizing”	2 h 34 min	2 h 34 min	2 h 14 min

several reasoning experiments that were conducted to find a good trade-off between reasoner runtime and complexity of ontology design.

6.2 Sequential vs. Parallel Assignment of Individuals to Classes

In a number of preparatory reasoning experiments, we made the following observations:

- During the whole runtime, the Pellet reasoner did not use more than approx. 4 Gb of main memory.
- The assignment of individuals to classes (signaled to the user by a small popup window announcing the reasoner activity “Realizing”) consumes almost all the Pellet runtime.
- Pellet makes use only of a single core, i.e., it runs strictly sequential during the “Realizing” phase.

With the goal to reduce the reasoner runtime significantly in mind, we concentrated on parallelization of the “Realizing” phase of the Pellet reasoner.

Pellet performs the classification of individuals in two loops that are nested: one loop over individuals, the other loop over classes, i.e., the inner loop is performed in full per iteration step of the outer loop. If we denote by m the number of classes and by n the number of individuals, the reasoner runtime thus depends at least on $m \cdot n$. Inserting the numbers given above for classes and individuals, $m \cdot n = 29,429,366$ holds for the SITC-4 ontology as proposed in this paper. In the extreme case, this means that the SITC-4 classification would cost approx. $30 \cdot 10^6$ possibly complex computations. It is neither possible to reduce the number of classes, nor the number of individuals.

The `pellet.properties` file holds a Boolean configuration variable named `REALIZE_INDIVIDUAL_AT_A_TIME` that lets the user choose which loop is the outer loop: When set to `TRUE`, the loop over individuals is the outer loop, when set to `FALSE`, the loop over classes is the outer loop.

For the Pellet option “Outer loop iterates over individuals” an inspection of the Pellet program code revealed that the iteration steps do not have mutual data dependencies. The absence of data dependencies allowed us to design a rather simple parallel version for the “Realizing” code, built upon the Java Executor Service. We inserted 12 additional lines of code and modified 3 lines. We did not

conduct any further optimizations like, e.g., adaption of the garbage collection frequency.

Unfortunately, there are mutual data dependencies in the Pellet code for iteration steps for the option “Outer loop iterates over classes”; to avoid error-prone re-programming, we did not try to implement a parallel version for this option. We discuss observed runtimes in the following chapter together with the effects of different ontology versions.

6.3 Implicit vs. Explicit Assignment of Individuals to Classes

We analyzed also if we could reduce reasoner runtimes by giving up the stated goal to let the reasoner do the complete job of assigning individuals to classes. We assumed that by introducing explicit `rdf:type` axioms per individual the classification task could be made easier for the reasoner. Obviously, this solution has the drawback that the reasoner’s potential to generate structural relationships between ontology elements is not very well exploited.

To study the effect of inserting explicit `rdf:type` axioms, we produced three subversions of SITC-4+, namely SITC-4+_{impl}, SITC-4+_{map}, and SITC-4+_{expl}:

- Subversion SITC-4+_{impl} (“implicitly typed individuals”) is identical to the initial SITC-4+, i.e. this subversion does not contain any `rdf:type` axioms neither for individuals representing SITC-4 basic headings, nor for CN Mappings.
- In subversion SITC-4+_{map} (“typed mapping individuals”) we introduced `rdf:type` axioms only for CN Mappings. As mentioned before, the SITC-4 ontology comprises 18,231 individuals of this kind.
- In subversion SITC-4+_{expl} (“explicitly typed individuals”) we introduced `rdf:type` axioms for all individuals.

Table 3 shows the measured mean runtimes, calculated from 3 runs per experiment. Pellet was running the extract subcommand with default parameters. The software environment comprised the Java SE Runtime Environment 1.7.0 provided by a Java 64-Bit Server VM on a Linux system, the hardware environment comprised an Intel Core i7 processor with 2.9 GHz clock rate and 8 cores. We discuss the shown figures as follows:

- For the SITC-4+_{impl} subversion, parallelization allowed us to reduce the reasoner runtime to approx. 80% of the observed reasoner runtime with sequential “Realizing” phase.
- The SITC-4+_{map} subversion does not reduce the observed runtimes, neither for the sequential, nor for the parallel “Realizing” phase.

- For the sequential “Realizing” phase, the SITC-4_{+expl} subversion drastically reduces the observed runtime to approx. 3% of the SITC-4_{+impl} subversion.
- The runtimes observed for the parallel “Realizing” phase are almost the same for all different SITC-4+ subversions.

Though a detailed analysis of the Pellet behavior is beyond the scope of this paper, we may assume that the Pellet option “Outer loop iterates over classes” is the conceptually better choice for our problem. We observed during several runs that the Pellet reasoner deals with the :CNMapping class in the first iteration step. This allows Pellet to get rid of the 18,231 CN Mappings already in this iteration step. The reasoner does not need to take these individuals into account in following iteration steps, because all classes are disjoint. After assigning the CN Mappings to the :CNMapping class, a membership in other classes is not possible for the CN Mappings. Per following iteration step, the reasoner must deal only with the remaining 2987 individuals. The observed runtimes suggest that assigning these 2987 individuals to the 1383 SITC-4 classes is very hard, when the reasoner must infer class membership from properties, and very easy, when class membership is explicitly declared by `rdf:type` axioms.

Parallel “Realizing” for the “Outer loop iterates over individuals” option of Pellet can compensate its inherent disadvantage by applying a large enough number of cores, but it remains the fact that the number of classes in the inner loop cannot be reduced in any iteration. The full number of approx. $30 \cdot 10^6$ complex computations must be carried out.

We think that for larger ontologies – populated with even more individuals – the Pellet “Outer loop iterates over classes” option may exhibit its advantage even more clearly. An effort to design a parallel version for the “Realizing” phase for this option seems very worthwhile.

7 Summary and Outlook

In this paper, we presented an approach to convert the Standard International Trade Classification into a semantically rich OWL-2 ontology. We discussed some design choices and reported from our experience with the Pellet reasoner. We believe that the outcome of our work – besides the actual SITC-4 ontology – is twofold: we learned that ontology development even for seemingly regular structures may become quite a complex engineering task, and that medium to large size ontologies require careful design that considers not only ontological structures but also reasoning and especially reasoning runtimes. The benefit of reasoning, though, is that verifiable structure

information can be added to data in an automatic procedure.

Both the developed ontology for SITC-4 and the XSL(T) style sheets will be published on the authors’ website (<http://comsys.informatik.uni-kiel.de>).

In our ongoing work, we concentrate on integrating trade data into the given ontology, and we extend our effort to the Combined Nomenclature.

References

- Caracciolo C, Heguiabehera J, Gangemi A, Baldassarre C, Keizer J, Taconet M (2012) Knowledge management at FAO: a case study on network of ontologies in fisheries. In: Suárez-Figueroa MC, Gómez-Pérez A, Motta E, Gangemi A (eds) *Ontology engineering in a networked world*. Springer, Heidelberg, pp 383–405
- Deutsche Zentralbibliothek für Wirtschaftswissenschaften. Digitale Reichsstatistik. <http://zbw.eu/de/ueber-uns/arbeitschwerpunkte/forschungsdatenmanagement/digitale-reichsstatistik/>. Accessed 2 Sep 2015
- Deutsche Zentralbibliothek für Wirtschaftswissenschaften. Standard Thesaurus Wirtschaft. <http://zbw.eu/stw/version/latest/about>. Accessed 2 Sep 2015
- Ding Y, Fensel D, Klein M, Omelayenko B, Schulten E (2004) The role of ontologies in ecommerce. In: Staab S, Studer R (eds) *Handbook of ontologies*. Springer, Heidelberg
- eCl@ss e.V. eCl@ss: classification and product description. <http://www.eclass.de/eclasscontent/index.html.en>. Accessed 2 Sep 2015
- European Commission. Combined nomenclature. <http://ec.europa.eu/eurostat/ramon/index.cfm>. Accessed 1 Sep 2015
- Eurostat (2015a) Correspondence table between standard international trade classification, rev. 4 and combined nomenclature 2015. http://ec.europa.eu/eurostat/ramon/other_documents/combined%20nomenclature/conversion_tables/CN2015_SITC4.zip. Accessed 1 Sep 2015
- Eurostat (2015b) Combined nomenclature ad hoc conversion tables. http://ec.europa.eu/eurostat/ramon/other_documents/combined%20nomenclature/conversion_tables/CN_SITC_2015.zip. Accessed 1 Sep 2015
- Food and Agriculture Organization of the United Nations (2002) CWP handbook of fishery statistical standards. Section R: fishery commodities classification. CWP data collection. In: FAO Fisheries and Aquaculture Department. Rome. <http://www.fao.org/fishery/cwp/handbook/R/en>. Accessed 26 Oct 2015
- Hepp M, de Bruijn J (2007) GenTax: a generic methodology for deriving OWL and RDF-S ontologies from hierarchical classifications, thesauri, and inconsistent taxonomies. In: Franconi E, Kifer M, May W (eds) *The semantic web: research and applications*, vol LNCS 4519. Springer, Heidelberg, pp 129–144
- International Organization for Standardization (2000) ISO 1087-1:2000, terminology work – vocabulary – part 1: theory and application
- International Organization for Standardization. ISO 25964 Information and documentation – thesauri and interoperability with other vocabularies. Citation from Wikipedia, https://en.wikipedia.org/wiki/ISO_25964. Accessed 2 Sep 2015
- Krieger H-U, Declerck T (2015) An OWL ontology for biographical knowledge. Representing time dependent factual knowledge. In: *Proceedings of the first conference on biographical data in a*

- digital world. <http://ceur-ws.org/Vol-1399/paper16.pdf>. Accessed 15 Oct 2015
- Krötzsch M, Simančík F, Horrocks I (2013) A description logic primer. [arXiv:1201.4089v3](https://arxiv.org/abs/1201.4089v3) [cs.AI] 3 Jun 2013. Accessed 20 Oct 2015
- Noy NF, McGuinness DL (2001) Ontology development 101: a guide to creating your first ontology. http://protege.stanford.edu/publications/ontology_development/ontology101.pdf. Accessed 15 Oct 2015
- Noy N, Rector A, Hayes P, Welty C (eds) (2006) Defining N-ary relations on the semantic web. W3C Working Group Note 12 April 2006, <http://www.w3.org/TR/swbp-n-aryRelations/>. Accessed 15 Oct 2015
- Staab S, Studer R (eds) (2009) Handbook on ontologies, 2nd edn. Springer, Heidelberg
- Statistisches Bundesamt (2006) Deutsche Übersetzung der Standard International Trade Classification, Revision 4, der Vereinten Nationen, Ausgabe 2006. https://www.destatis.de/DE/Methoden/Klassifikationen/Aussenhandel/DeutscheFassungSITC.pdf?__blob=publicationFile. Accessed 2 Sep 2015
- Stolz A, Rodriguez-Castro B, Radinger A, Hepp M (2014) PCS2OWL: a generic approach for deriving web ontologies from product classification systems. In: Presutti V, d'Amato C, Gandon F, d'Aquin M, Staab S, Tordai A (eds) The semantic web: trends and Challenges, 11th international conference ESWC 2014. LNCS, vol 8465. Anissaras, Crete, Greece, pp 644–658
- Tolksdorf R, Bizer Ch, Eckstein R, Heese R (2003) Business to consumer markets on the semantic web. In: Meersman R, Tari Z (eds) On the move to meaningful internet systems, OTM 2003-WMS, Springer LNCS 2889, pp 816–828
- UDC Consortium. Universal decimal classification. <http://www.udcc.org/>. Accessed 2 Sep 2015
- United Nations Statistics Division (2010) Standard international trade classification, revision 4. <http://unstats.un.org/unsd/trade/sitcrev4.htm>. Accessed 1 Sep 2015
- United Nations Statistics Division (2015) Standard international trade classification (SITC) revision 4. <http://unstats.un.org/unsd/pubs/gesgrid.asp?id=104>. Accessed 1 Sep 2015
- UNSPSC (2014) UN development programme: United Nations standard products and services code. <http://www.unspsc.org/>. Accessed 2 Sep 2015
- W3C (2012) SKOS simple knowledge organization system. <http://www.w3.org/2004/02/skos/>. Accessed 2 Sep 2015
- W3C. ConverterToRdf. <http://www.w3.org/wiki/ConverterToRdf>. Accessed 20 Oct 2015
- World Customs Organization. What is the harmonized system? <http://www.wcoomd.org/en/topics/nomenclature/overview/what-is-the-harmonized-system.aspx>. Accessed 1 Sep 2015