

# **Using Semantic Web Tools for Context Interchange**

**Mihai Lupu  
Stuart E. Madnick**

**Working Paper CISL# 2007-07**

**September 2007**

Composite Information Systems Laboratory (CISL)  
Sloan School of Management, Room E53-320  
Massachusetts Institute of Technology  
Cambridge, MA 02142

# Using Semantic Web tools for COnText INterchange

Mihai Lupu  
Singapore-MIT Alliance  
National University of  
Singapore  
mihailup@comp.nus.edu.sg

Stuart Madnick  
Sloan School of Management  
Massachusetts Institute of  
Technology  
smadnick@mit.edu

## ABSTRACT

The COnText INterchange Strategy (COIN) is an approach to solving the problem of interoperability of semantically heterogeneous data sources through context mediation. The existing implementation of COIN uses its own notation and syntax for representing ontologies. More recently, the OWL Web Ontology Language is becoming established as the W3C recommended ontology language. A bridge is needed between these two areas and an explanation on how each of the two approaches can learn from each other. We propose the use of the COIN strategy to solve context disparity and ontology interoperability problems in the emerging Semantic Web both at the ontology level and at the data level. In this work we showcase how the problems that arise from context-dependant representation of facts can be mitigated by Semantic Web techniques, as tools of the conceptual framework developed over 15 years of COIN research.

## 1. INTRODUCTION

Making computers understand humans is, generously put, a hard task. One of the main reasons for which this is such a hard task is because even humans cannot understand each other all the time. Even if we all spoke the same language, there still exist plenty of opportunities for misunderstanding. An excellent example is that of measure units. Again, we don't even have to go across different names to find differences: in the US, a *gallon* (the so-called Winchester gallon) is approximately 3785 ml while in the UK, the "same" *gallon* is 4546 ml, almost 1 liter more. So when we find a piece of information in a database on cars, for instance, and we learn that a particular model has a fuel tank capacity of 15, how much gas can we actually fit inside, and, consequently, for how long can we drive without stopping at a gas station?

The answer to the previous problem comes easy if we know where we got the data from: if the information was from the US, we know we can fit inside 56.78 liters of gas, while if it comes from the UK, it is 68.19 - a difference of about 11 liters, with which a car might go for another 100 miles (or

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permission from the publisher, ACM.

VLDB '07, September 23-28, 2007, Vienna, Austria.

Copyright 2007 VLDB Endowment, ACM 978-1-59593-649-3/07/09.

161 km if the driver is not American or British).

Many more such examples exist (see [10] for a particularly costly one) and the reason for which they persist is mainly because it is hard to change the schema of relational databases that do not include the units of their measurements simply because when they were designed, they were designed for a single context, where everybody would know what the units are. With globalization, off-shoring, outsourcing and all the other traits of the modern economical environment, those assumptions become an obstacle to conducting efficient business processes.

Even in the context of a purely-semantic web application, such as the Potluck<sup>1</sup> project developed at the Computer Science and AI Laboratory at MIT, contextual information is not explicitly approached. The user is allowed to mash-up together information from different sites, but it is not taken into account the fact that those different data sources may have different assumptions about an entire array of concepts. This paper shows how the COIN strategy can be implemented in this new environment and how it can contribute to it.

### 1.1 Approach overview

Semantic web tools rely heavily on mathematical logic to perform inferences. The result of this, in combination with our desire to maintain a 100% pure logic approach, is that *facts* cannot be deleted or modified. For instance, even if we define a relation *hasName* to be of functional type (i.e. have a unique object for each subject), it is still legal to have two entries with different objects, such as (*Location1*,*hasName*, 'London') and (*Location1*,*hasName*, 'Londres'), the conclusion of which will be that the names "London" and "Londres" denote the same location. This has both advantages and disadvantages which we will not discuss here, but refer the reader to a wealth of literature on mathematical logic. Instead, we focus on how we model context in this framework.

Taking a simple *Weather* example [7] containing locations and their respective temperatures, Listing 1 shows how we might represent different values of the same temperature using prefixes of relations.

**Listing 1: A possible solution to representing context, by adding an additional *hasValue* relation to the Temperature object. It is simple and intuitive of the fact that we are dealing indeed with the same temperature.**

```
<?xml version="1.0" ?>
```

1

<sup>1</sup><http://dfhuynh.csail.mit.edu:6666/potluck/>

```

<rdf:RDF                                2
  <UK:Temperature rdf:ID="temp1">        3
    <UK:hasValue rdf:datatype="[#int]">  4
      10</UK:hasValue>                   5
    <USA:hasValue rdf:datatype='[#int]'>  6
      50</USA:hasValue>                   7
  </UK:Temperature>                       8
  <UK:Location rdf:ID="London">          9
    <UK:hasName rdf:datatype="[#string]"> 10
      London, UK</UK:hasName>            11
    <UK:hasTemperature rdf:resource="#temp1"/> 12
  </UK:Location>                          13
</rdf:RDF>                                14

```

The way to add this new relation is by defining a SWRL [5] rule such as the one in Listing 2 and then querying the results with a query in SPARQL [12] as in Listing 3.

**Listing 2: SWRL rule that generates a value in the USA context**

```

UK:hasValue(?temp, ?tempValue) ^        1
UK:hasTemperature(?location, ?temp) ^    2
swrlb:multiply(?product,?tempValue,1.8) ^ 3
swrlb:add(?sum, ?product, 32) →          4
USA:hasValue(?temp, ?sum)                5

```

**Listing 3: SPARQL query that retrieves the value in the US context**

```

SELECT ?tempValue                          1
WHERE { ?1 UK:hasTemperature ?temp .       2
        ?1 UK:hasName "London" .          3
        ?temp USA:hasValue ?tempValue     4
      }                                     5

```

This all seems very intuitive. As always, the problems lie in the details: How do we determine a conflict of contexts? How do we identify the correct rule to be applied? How should the data and rules be organized into files? Do we apply the rule to the entire dataset thus generating massive amounts of new data, or should we just apply it to the subset being queried?

We answer all of these questions in the following sections. Among these, Section 4 introduces the basic representation of context using the Ontology Web Language (OWL [9]) and Section 5 presents our conflict identification and resolution method. A more detailed discussion, is available in [7].

## 2. BACKGROUND AND RELATED WORK

### 2.1 COntext INterchange

The idea of the COntext INterchange [4] system is to reuse massive amounts of data that already exist but that are incomplete due to design assumptions that eliminated constants from the dataset. When two or more such datasets are put together, or queried together, what used to be constants in each of them become variables in the aggregated dataset and consequently needs to be added back in the data. This is in most cases unfeasible due to the rigidity of the data structures or simply due to the fact that the end user has no control over the repository where the data exists.

The core of the COntext INterchange approach is a context mediator that rewrites queries coming from a user context into a context-sensitive mediated query that addresses the differences in meaning between the receiver and the sources. Conceptually, the context mediator is structured around a *domain model* that consists of *semantic types*, *attributes* and *modifiers*.

A semantic type is, as the name indicates, a conceptual entity. For instance, in the *Weather* example of Section 3, the column **temperature** has no meaning by itself until it is associated with a semantic type *Temperature*. The coincidence in names is just because humans created both entities, but it should be clear to the reader that the column could very well have been named **temp** and the semantic type *ST2842*. The important difference between semantic types and the columns with which they are associated is that the semantic types come enriched with semantics and attributes. In this simple example, *Temperature* has only one attribute, *value*.

In turns, an attribute may come endowed with a modifier. Again, using the *Weather* example, we can imagine that the *value* attribute of the *Temperature* semantic type has a modifier *unit*. This is called a *modifier* because it changes the meaning of the attribute to which it refers - it modifies it. A *modifier value* is in this case, *Celsius (C)* or *Fahrenheit (F)*.

COIN uses this architecture to automatically determine differences in contexts and resolve queries in a way that is easy to interpret correctly by the user, even if the data is expressed in a different context. Existing application include financial reporting and analysis, airfare and car-rental aggregators, etc. [3].

### 2.2 A glance at the Semantic Web

Though the COIN methodology precedes the Semantic Web, and despite the many similarities in objectives and motivation, the two have developed mostly independently. The wide spectrum of tools that have been proposed by different research groups to achieve the targets of the original paper by Berners-Lee et al. [2] make a quick but complete summary virtually impossible. In this section we just look at the few tools that we identify to be “best”, both in terms of the appropriateness for our own purposes, and also in terms of their acceptance and popularity within the Semantic Web community.

Clearly, one of the pillars of current Semantic Web research is the Web Ontology Language (OWL [9]). To query the data stored in OWL format, one could map it back to a relational database and query it with SQL or use a “native” query language such as SPARQL [12]. For most purposes, SPARQL can be translated back to SQL, but the advantage of it lies in being able to query directly the RDF graph that underlies any ontology. It has the status of *Working Draft* of the W3C since October 2006.

After representation and query languages, the Semantic Web framework requires a *rule language* to make inferences on the existing data, thus enabling the creation of the smart agents described in the original Berners-Lee paper. Though SWRL [5] has gained most attention in the past few years, the language has not yet been standardized by the W3C and many different implementations exist, that rarely support the full specification, mainly because in that case the reasoning becomes undecidable. One of the most popular implementation is SWRLTab<sup>2</sup> - an extension to the Protégé framework<sup>3</sup>, that uses mainly the Jess<sup>4</sup> inference engine (though it could use other engines too).

<sup>2</sup><http://protege.cim3.net/cgi-bin/wiki.pl?SWRLTab>

<sup>3</sup><http://protege.stanford.edu/>

<sup>4</sup><http://herzberg.ca.sandia.gov/jess/>

### 3. FROM TABLES TO INFORMATION

The first step towards making the data understandable by different agents<sup>5</sup> performing their activities in different contexts is to understand the fact that we are only dealing with representations of concepts and facts. As we exemplified before,  $15 = 56.78 = 12.49$  if one is gallons (US), one is liters and one is gallons (UK). Consequently, it makes more sense to have an abstract concept representing this volume and attach to it the knowledge that it may be expressed in different ways.

A tempting way of moving information out of the restrictive relational database is to encode it using XML. Using a naïve method implemented in most database systems, this would result in, literally, a data dump. For instance, a simple table containing locations and temperatures (Table 1) can be expressed as in Listing 4.

Using XML does not solve our problem. As discussed in [8], XML is not a silver bullet - it is just another way to express the data. It only provides a more flexible way, allowing us to add more meaning to it. A simple “data dump” from the relational database is not enough for two reasons: First, as we see in Listing 4, the file mixes together the structure of the data with the data itself. Conceptually, these are different and should be represented as such. Second, the data itself is stored as if to preserve the physical appearance of the table (i.e. a sequence of rows, each with a few columns) rather than to preserve its underlying meaning. It is thus clear that a different approach is needed.

Location	Temperature
london	10
paris	12

Table 1: Sample relational table

Listing 4: XML representation of relational database

```

1 <?xml version="1.0"?>
2 <mysqldump xmlns:xsi="..."XMLSchema-instance">
3 <database name="test">
4   <table_structure name="weatherUK">
5     <field Field="location" Type="varchar(20)" />
6     <field Field="temperature" Type="int(11)" />
7   </table_structure>
8   <table_data name="weatherUK">
9     <row>
10      <field name="location">london</field>
11      <field name="temperature">10</field>
12    </row>
13    [...]
14  </table_data>
15 </database>
16 </mysqldump>

```

There exist attempts to extract ontological information from relational tables [1, 6]. What we want here is not nearly as ambitious as the works mentioned before. For our purpose, we don't necessarily need to infer a full scale ontology from the data, but simply to express things that are the same as being the same and things that are different as being different. It sounds simple for a human being, but computers have serious difficulties in performing even this simple task.

The first thing we want to do is separate the structure of the representation from the data itself. Listing 5 shows

<sup>5</sup>we prefer the term 'agents' to show that they can be either human end-users or other computer systems

how we can define an ontological structure to organize the data in the table. We use the term “ontological” simply because we use the ontology specification language, OWL, but one should not imagine a complex theory behind it: in this listing we simply state that we deal with two concepts (*Location* and *Temperature*) who are connected by a relationship *hasTemperature*. The difference between this approach and the simple XML dump is that here *location* and *temperature* are regarded as concepts, rather than fields in a table. It is a subtle, but essential difference. Here, a particular instance of the *Location* class has a name, but is separate from its name. This distinction will allow us later to specify that *London* is the same location with *Londres* and that 10 is the same temperature with 50 (one in Celsius and one in Fahrenheit degrees). This way, in the data file shown in Listing 6 we can define the abstract temperature *temp1* and give it a value and then define the abstract location *London* and give it a name and associate it with the abstract temperature value. In these listings, an *ObjectProperty* relates two instances of two classes, while a *DatatypeProperty* relates the instance of a class to a pre-defined type (integer, string, etc.).

Listing 5: legacyUK.owl:Ontology structure for the information in the relational database

```

1 <?xml version="1.0"?>
2 <rdf:RDF [...]
3   xml:base="legacyUK.owl">
4     <owl:Ontology rdf:about="" />
5     <owl:Class rdf:ID="Location" />
6     <owl:Class rdf:ID="Temperature" />
7     <owl:ObjectProperty rdf:ID="hasTemperature">
8       <rdfs:domain rdf:resource="#Location" />
9       <rdfs:range rdf:resource="#Temperature" />
10    </owl:ObjectProperty>
11    <owl:DatatypeProperty rdf:ID="hasValue">
12      <rdfs:range rdf:resource="[#int]" />
13      <rdfs:domain rdf:resource="#Temperature" />
14    </owl:DatatypeProperty>
15    <owl:FunctionalProperty rdf:ID="hasName">
16      <rdfs:range rdf:resource="[#string]" />
17      <rdfs:domain rdf:resource="#Location" />
18      <rdfs:type rdf:resource="#DatatypeProperty" />
19    </owl:FunctionalProperty>
20  </rdf:RDF>

```

Listing 6: legacyUKdata.owl:Data represented using the ontological structure

```

1 <?xml version="1.0"?>
2 <rdf:RDF [...]
3   xmlns:UK="legacyUK.owl#"
4   xmlns:cntxts="cntxts.owl#"
5   xml:base="legacyUKdata.owl">
6 <owl:Ontology rdf:about="">
7 <owl:imports>
8 <rdf:Description rdf:about="legacyUK.owl">
9 </rdf:Description>
10 </owl:imports>
11 </owl:Ontology>
12 <UK:Temperature rdf:ID="temp1">
13   <UK:hasValue [...]>10</UK:hasValue>
14 </UK:Temperature>
15 <UK:Location rdf:ID='London'>
16 <UK:hasName [...]>London, UK</UK:hasName>
17   <UK:hasTemperature rdf:resource='temp1' />
18 </UK:Location>[...]
19 </rdf:RDF>

```

Listings 5 and 6 show the kind of input our system considers as *source*. With respect to the amount of reasoning

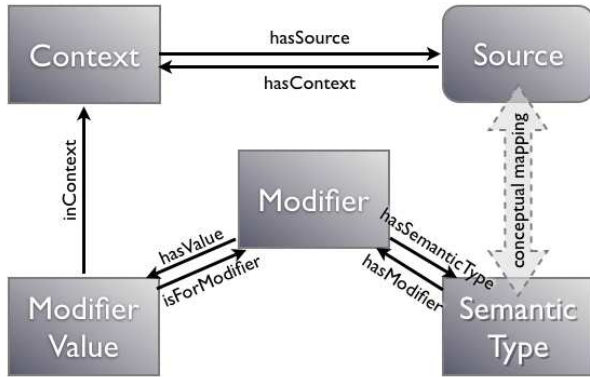


Figure 1: Our model: apart from the source, three components are needed, and the *Modifier* is the binding concept

necessary at this point, our requirements are quite low since the machine needs not understand the concepts, but merely identify them as concepts rather than rows or columns in a table. The translation from the relational-model representation to our ontological representation is easily done automatically using one of the several available transformation languages such as XSLT, FleXML or HaXml.

#### 4. SEPARATING CONTEXT FROM DATA REPRESENTATION

In order to be able to do the things outlined in Section 1.1 we first need to establish a way to represent context. The flexibility of the RDF and OWL languages allow for such a variety of architectures to be defined, that one of the problems we faced was focusing on one in particular, one that provides, in our opinion, the best solution for future extensions.

Initially, we had reduced the possibilities to three models, as detailed in [7]. Finally, the chosen model considers the semantic types to be independent of both the context and the source. In fact, we should imagine these semantic types as defined in an external ontology. This method provides the most independence between the different concepts. We will be using three files to express context (rectangles in Figure 1), in addition to the two files that represent the data (the ontology and the instances represented together as *Source* in Figure 1).

First, Listing 7 defines what a context is: a set of modifiers attached to a semantic type (some items present in the file will be explained in the next sections).

Listing 7: cntxtDefs.owl: the context definition

```
<?xml version="1.0"?> 1
<rdf:RDF 2
  xmlns:p1="cntxtDefs.owl#" 3
  xml:base="cntxtDefs.owl"> 4
  <owl:Ontology rdf:about=""/> 5
  <owl:Class rdf:ID="Context"/> 6
  <owl:Class rdf:ID="Query"/> 7
  <owl:Class rdf:ID="TriggeredRules"/> 8
  <owl:Class rdf:ID="SemanticType"/> 9
  <owl:Class rdf:ID="Modifier"/> 10
  <owl:Class rdf:ID="ModifierValue"/> 11
  <owl:ObjectProperty rdf:ID="hasModifiers"> 12
    <rdfs:range rdf:resource="#Modifier"/> 13
    <rdfs:domain rdf:resource="#SemanticType"/> 14
  </owl:ObjectProperty> 15
  <owl:ObjectProperty rdf:ID="isSemanticType"> 16
```

```
<rdfs:range rdf:resource="#SemanticType"/> 17
</owl:ObjectProperty> 18
<owl:DatatypeProperty rdf:ID="hasValue"> 19
  <rdfs:range rdf:resource="[#string]"/> 20
  <rdfs:domain rdf:resource="#ModifierValue"/> 21
</owl:DatatypeProperty> 22
<owl:ObjectProperty rdf:ID="hasContext"> 23
  <owl:ObjectProperty rdf:ID="inContext"> 24
    <rdfs:range rdf:resource="#Context"/> 25
    <rdfs:domain rdf:resource="#ModifierValue"/> 26
  </owl:ObjectProperty> 27
  <owl:ObjectProperty rdf:ID="isForModifier"> 28
    <rdfs:range rdf:resource="#Modifier"/> 29
  </owl:ObjectProperty> 30
  <owl:DatatypeProperty rdf:ID="ruleName"> 31
    <rdfs:domain rdf:resource="#TriggeredRules"/> 32
    <rdfs:range rdf:resource="[#string]"/> 33
  </owl:DatatypeProperty>[...] 34
</rdf:RDF> 35
```

The following listing (Listing 8) contains the file that represents the system's *understanding of the world*: a list of concepts, what properties they have (modifiers) and how they relate to each other.

Listing 8: cntxtOntology.owl: Semantic types definitions (including the modifiers they accept)

```
<?xml version="1.0"?> 1
<rdf:RDF 2
  xmlns="cntxtOntology.owl#" 3
  xmlns:cntxtDef="cntxtDefs.owl#" 4
  xml:base="cntxtOntology.owl"> 5
  <owl:Ontology rdf:about=""> 6
    <owl:imports rdf:resource="cntxtDefs.owl"/> 7
  </owl:Ontology> 8
  <cntxtDef:SemanticType rdf:ID="Temperature"> 9
    <cntxtDef:hasModifiers> 10
      <cntxtDef:Modifier rdf:ID="TemperatureUnit"/> 11
    </cntxtDef:hasModifiers> 12
  </cntxtDef:SemanticType> 13
  <cntxtDef:SemanticType rdf:ID="Location"/> 14
</rdf:RDF> 15
```

Finally, Listing 9 shows the instance of a context: all, or just a subset of modifiers, are assigned values in this file.

Listing 9: UKCntxt.owl: Context instance file

```
<?xml version="1.0"?> 1
<rdf:RDF 2
  xmlns="UKcntxt.owl#" 3
  xmlns:cntxtOntology="cntxtOntology.owl#" 4
  xmlns:cntxtDef="cntxtDefs.owl#" 5
  xml:base="UKcntxt.owl"> 6
  <owl:Ontology rdf:about=""> 7
    <owl:imports rdf:resource="cntxtOntology.owl"/> 8
  </owl:Ontology> 9
  <cntxtDef:ModifierValue rdf:ID="C"> 10
    <cntxtDef:inContext rdf:resource="#UKContext"/> 11
    <cntxtDef:isForModifier rdf:resource=" 12
      cntxtOntology.owl#TemperatureUnit"/> 13
    <cntxtDef:hasValue rdf:datatype="[#string] 14
      >C/</cntxtDef:hasValue> 15
  </cntxtDef:ModifierValue> 16
</rdf:RDF> 17
```

One problem with the structure presented above is that a basic query cannot differentiate between the two values of a modifier, even if they were defined in different files. This issue can be overcome using the *from* named construct in the SPARQL query language, as we show in [7].

## 4.1 The mediator file

To make the system work, one file needs to import all these bits and pieces together and build the construct of the COntext INtegration strategy. We call this the *mediator* file. Listing 10 shows an extract of its contents, in particular the `import` statements and the way we define the context of a source. In this example, the source is just the Temperature entity `temp1` defined in the `UKInstances.owl` file. However, it can be anything else: an entire file, a class or just an instance as in this case.

**Listing 10: The mediator file puts together all the different pieces of the architecture and defines particular contexts of the sources**

```

<owl:Ontology rdf:about=""> 1
  <owl:imports rdf:resource="USAcntxt.owl"/> 2
  <owl:imports rdf:resource="UKcntxt.owl"/> 3
  <owl:imports rdf:resource="UKinstances.owl"/> 4
</owl:Ontology> 5
<rdf:Description rdf:about="UKinstances.owl#temp1"> 6
  <j.0:isSemanticType rdf:resource= 7
    "cntxtOntology.owl#Temperature"/> 8
  <j.0:hasCntxt rdf:resource="UKcntxt.owl#"/> 9
</rdf:Description> 10

```

Now we can, for instance, identify the context of a source using a query similar to the following:

**Listing 11: Query to identify the context of a source**

```

prefix cntxtDefs:<cntxtDefs.owl#> 1
prefix mediator:<all.owl#> 2
select ?data ?context 3
from <all.owl> 4
where {?data cntxtDefs:hasCntxt ?context} 5

```

## 5. CONTEXT CONFLICT IDENTIFICATION AND RESOLUTION

In the previous sections we have explained how a user might query the data to find out what is the appropriate context it refers to. In this section we will show how we do this automatically for the purpose of context conflict determination and how this determination will trigger the necessary conversion rules. We will continue to use the example of locations and temperatures presented throughout this work.

The approach we follow in this work is a two-step approach: first, we need to determine the need for a conversion (i.e. determine the existence of two different contexts) and then apply the corresponding rule.

These two phases are implemented in two sets of rules: first a *trigger rule* analyses the data and the query to identify potential conflicts. If one such conflict is identified, a flag is raised, to announce the necessity of the application of a conversion rule. We will describe the implementation of this flag shortly. Upon assertion of the trigger flag, the corresponding rule will automatically be triggered and context mediation will take place by addition of new data to the dataset.

*The trigger rule.* Listing 12 shows the exact rule used to determine conflicts between any two attributes of the same type. Lines 1-4 identify the difference between the contexts of the data and the query. Then, lines 5-6 identify the semantic type and modifier, lines 7-9 identify the value of the modifier in the context of the dataset while lines 10-12 do the same for the value of the modifier in the query context.

Finally, the test is performed in line 13 and if the two values are different, a trigger is built and added to the dataset. In this implementation, the attribute itself is linked by a `hasCntxt` relation to a particular context. In other situations, such a relation may only be defined for the entire dataset, rather than for individual attributes. This is not a problem, as a SWRL rule can extend the `hasCntxt` rule from a class to its components.

**Listing 12: Rule for the determination of context conflict**

```

UKdefs:hasValue(?attribute, ?attributeValue) ^ 1
cntxtDefs:hasCntxt(?attribute, ?dataContext) ^ 2
cntxtDefs:hasCntxt(Query_1, ?queryContext) ^ 3
differentFrom(?dataContext, ?queryContext) ^ 4
cntxtDefs:isSemanticType(?temp, ?semType) ^ 5
cntxtDefs:hasModifiers(?semType, ?modifier) ^ 6
cntxtDefs:isForModifier(?modVal, ?modifier) ^ 7
cntxtDefs:hasValue(?modVal, ?dataModVal) ^ 8
cntxtDefs:inContext(?modVal, ?dataCntxt) ^ 9
cntxtDefs:isForModifier(?modVal1, ?modifier) ^ 10
cntxtDefs:hasValue(?modVal1, ?queryModVal) ^ 11
cntxtDefs:inContext(?modVal1, ?queryCntxt) ^ 12
differentFrom(?dataModVal, ?queryModVal) ^ 13
swrlb:stringConcat(?trigger1, "-to-", ?queryModVal) ^ 14
swrlb:stringConcat(?trigger, ?dataModVal, 15
?queryModVal) 16
→ cntxtDefs:ruleName(TriggeredRules1, ?trigger) 17

```

*The conversion rule.* The actual conversion rule that transforms degrees in Celsius to degrees in Fahrenheit is shown in Listing 13. Line 1 checks the existence of the triggered flag and, if this condition is satisfied, it performs the necessary mathematical conversion functions (lines 2-5) and asserts the new value in line 6.

The result of applying both rules on the knowledge base is shown in Figure 2. We can see that two new facts have been asserted: first, the trigger rule has discovered the context conflict and, second, upon assertion of the conflict, the conversion rule has been triggered to compute the new value.

**Listing 13: Conversion rule**

```

cntxtDefs:ruleName(TriggeredRules1, "C-to-F") ^ 1
UKdefs:hasValue(?temp, ?tempValue) ^ 2
swrlb:multiply(?temp1, ?tempValue, 9) ^ 3
swrlb:divide(?temp2, ?temp1, 5) ^ 4
swrlb:add(?newValue, ?temp2, 32) 5
→ USAdefs:hasValue(?temp, ?newValue) 6

```

*Query alteration.* The simple way in which we have defined the conversion rule (Listing 13) allows us to re-write the query in an equally simple manner. Listings 14 and 15 show the original and, respectively, the new query for obtaining the value of temperature in London. As it can be observed, the only difference is the prefix of the `hasValue` relationship.

**Listing 14: Original query**

```

SELECT ?tempValue 1
WHERE { ?loc UKdefs:hasTemperature ?temp. 2
        ?loc UKdefs:hasName "London" . 3
        ?temp UKdefs:hasValue ?tempValue} 4

```

**Listing 15: New query for the USA context**

```

SELECT ?tempValue 1
WHERE { ?loc UKdefs:hasTemperature ?temp. 2

```

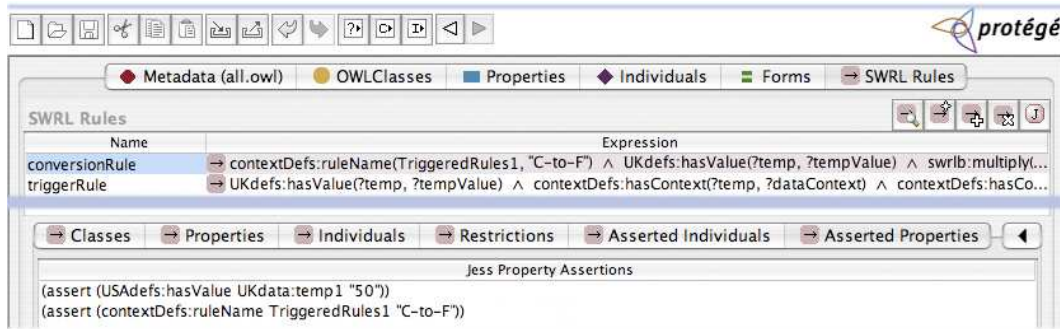


Figure 2: With two rules in the knowledge base, the conversion between temperature units has been performed automatically. In the lower half of the image we can see that the new temperature value (50F) has been correctly asserted

?loc UKdefs:hasName "London" .	3
?temp USAdefs:hasValue ?tempValue}	4

Despite its ease of use and implementation, the current method is not perfect. The relation `USAdefs:hasValue` is not directly linked to the USA context. Formally, it has no link to any specific context. Though this approach can be implemented programatically, our future work, described in the next section, aims towards an ever closer integration with the semantic web tools.

## 6. FUTURE WORK

Our work so far has shown how we can approach the problem of context interchange using the COIN strategy via the tools of the Semantic web. To fully achieve all the features that are currently available in COIN there are still steps ahead, some of which we describe in this section.

The solution presented in the previous section relies on external programming languages to transform a query such that it returns the result in a different context. A better solution would be to have a new tertiary relation, similar to the one that defined the value of a modifier in a particular context. This new relation, which we call `hasCntxtValue` links together an attribute, a value and a context. As we have seen, SWRL can only express binary relations directly, so the only way to implement this relation is to define it as an owl:Class with three binary relations. Now, the conversion rule needs to infer the new tertiary relation that links the attribute to the new value in the new context. In [7] we have showed how such a rule can be created following the Semantic Web best practices [11].

The difficulty in inferring this relation in the conversion rule is that a new instance has to be generated: a new individual of the `hasCntxtValue` type that would link the three components (attribute, value and context). Unfortunately, the current standard SWRL specification does not provide means to instantiate classes, thus making this solution temporarily unfeasible.

## 7. CONCLUSION

In this work we describe how the COntext INterchange strategy can be implemented using the Semantic Web tools, in particular using OWL, SWRL and SPARQL. We acknowledge the existence of massive amounts of data in relational databases that lack all the necessary data required for users other than the original designers of the database and describe how the information present in these databases can

be “elevated” to a knowledge base. Subsequently, we show how to structure information pertaining to the context of the data - how to model the definitions of *semantic type*, *modifier* and *modifier value*. Using these models we show how the necessary conversions of the data values can be made by using a two-step process involving pairs of *trigger* and *conversion* rules.

## 8. REFERENCES

- [1] I. Astrova. *The Semantic Web: Research and Applications*, chapter Reverse Engineering of Relational Databases to Ontologies. Springer Berlin / Heidelberg, 2004.
- [2] T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, 2001.
- [3] A. Firat. *Information Integration using Contextual Knowledge and Ontology Merging*. PhD thesis, Sloan Business School, MIT, 2003.
- [4] C. H. Goh, S. Bressan, S. Madnick, and M. Siegel. Context interchange: New features and formalisms for the intelligent integration of information. *ACM TIS*, 17(3), 1999.
- [5] I. Horrocks, P. Patel-Schneider, H. Boley, S. Tabet, B. Groszof, and M. Dean. SWRL: A semantic web rule language combining OWL and RuleML. <http://www.w3.org/Submission/SWRL/>, 2004.
- [6] N. Lammari and E. Metais. Building and maintaining ontologies: a set of algorithms. *NLDB*, 48(2), 2004.
- [7] M. Lupu and S. Madnick. On the feasibility of using semantic web tools for context interchange. <http://www.comp.nus.edu.sg/~g0404405/coinsw.WP.pdf>, 2007.
- [8] S. Madnick. The misguided silver bullet: What XML will and will not do to help information integration. In *Procs. of the iiWAS*, 2001.
- [9] D. McGuinness and F. van Harmelen. Owl web ontology language overview. <http://www.w3.org/TR/owl-features/>, 2004.
- [10] NASA. Mars climate orbiter failure causes. <http://mars.jpl.nasa.gov/msp98/news/mco990930.html>.
- [11] N. Noy, A. Rector, P. Hayes, and C. Welty. Defining n-ary relations on the semantic web. <http://www.w3.org/TR/swbp-n-aryRelations/>, 2006.
- [12] E. Prud'hommeaux and A. Seaborne. Sparql query language for rdf. <http://www.w3.org/TR/rdf-sparql-query/>, 2007.