

# Querying Semantic Web Resources Using TRIPLE Views

Zoltán Miklós<sup>1</sup>, Gustaf Neumann<sup>1</sup>, Uwe Zdun<sup>1</sup>, and Michael Sintek<sup>2</sup>

<sup>1</sup> Vienna University of Economics, Department for Information Systems  
{zoltan.miklos, gustaf.neumann, uwe.zdun}@wu-wien.ac.at

<sup>2</sup> DFKI GmbH, Kaiserslautern, sintek@dfki.de

**Abstract.** Resources on the Semantic Web are described by metadata related to some formal or informal ontology. It is a common situation that a casual user does not know domain ontology in detail. This makes it difficult to formulate queries in this ontology to find the relevant resources. Users consider the resources in their specific context, so the most straightforward solution is to formulate queries in an ontology that corresponds to a user-specific view. We present an approach based on multiple views, expressed in simple ontologies. This allows a user to query heterogeneous data repositories in terms of multiple, relatively simple view ontologies. We present how ontology developers can define such views on ontologies and the corresponding mapping rules. These ontologies are represented in Semantic Web ontology languages, like RDFS, DAML+OIL or OWL. We present our approach with examples from the e-learning domain using the Semantic Web query and transformation language TRIPLE.

## 1 Introduction

One of the main motivations for the development of the Semantic Web is that the expansion of the World Wide Web makes it increasingly difficult for users to obtain relevant information. On the Semantic Web resources are annotated with metadata related to some ontology. Ontologies are formal explicit specifications of a shared conceptualization. Ontologies play a central role on the Semantic Web because they represent the relations among semantic information, hence all query or reasoning services have to be based on ontologies. Several ontology languages have been proposed and investigated recently, such as RDFS [11], DAML+OIL [12], OWL[13].

There are several problems which make the querying for Semantic Web resources difficult:

- *User Perspective:* Users who are interested to find resources formulate queries related to some ontology. The user's (or their agent's) view of a domain is often different than a domain expert's view. On the other hand, users are expected to be familiar with the domain ontology to efficiently formulate queries. Users could avoid to learn the domain ontology if they could formulate their queries in an ontology that corresponds to their domain-specific view.

- *Application Integration:* In the current state of the Semantic Web, ontologies are developed from scratch, therefore many ontologies describing the same domain exist. It is also widely believed that in the future development of the Semantic Web this ontology heterogeneity will remain and multiple ontologies for any particular domain will coexist. Applications face the problem to obtain information from sources which are described by different, independently developed ontologies.
- *Performance Overhead:* In many Semantic Web domains large or very large data sets exist. Queries can produce a considerable performance overhead. Problem-specific views of the resources could potentially minimize this problem.
- *Lack of formal definitions:* In the domain of e-learning, brokerage platforms (such as UNIVERSAL [18]) and RDF-based peer-to-peer networks (such as Edutella [19]) were developed, which act as a common mediators to access multiple data sources. Although during the last decade various standardized vocabularies for metadata were developed (e.g. LOM, SCORM), they are mostly informally defined and do not allow deeper reasoning. Furthermore, it is currently rather hopeless to define in this domain a large common ontology that covers aspects like the teaching domain, didactic goals, technical aspects of the delivery system, and so on.

In this paper, we propose an approach to overcome the problem of heterogeneous ontologies. Rather than defining a large single ontology, we propose to distinguish one or more “source ontologies” and a number of “target ontologies”. The source ontologies describe the resource instances and are usually rather complex. The target ontologies individually cover certain aspects of the application domain. A user who wants to query the heterogeneous data sources can formulate queries in terms of (multiple) relatively simple target ontologies. A view of the resources, described in the source ontologies, is created. In the view the resources are expressed with the target ontologies.

For our implementation we used the Semantic Web query and transformation language TRIPLE [2]. TRIPLE does not limit ourselves to use the same ontology language for the used ontologies. We specify the view mappings with description logic. We demonstrate that a specific class of mappings which is highly relevant in practical situations can be expressed in TRIPLE.

We present our approach firstly using an introductory example from the e-learning domain in Section 2. In Section 3 we describe our solution. In Section 4 we give a short overview of the language TRIPLE. Section 5 describes our approach using the previously introduced scenario in TRIPLE. In Section 6 we give a description Logic specification for our mappings. Section 7 provides related work and Section 8 concludes the paper with our plans for future research.

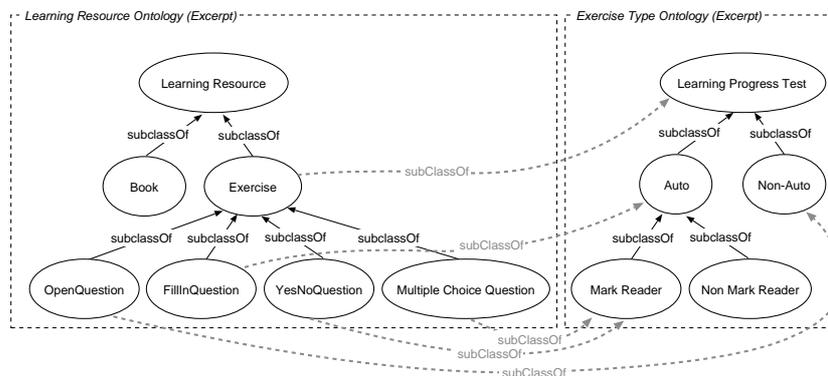
## 2 Example: E-learning Scenario

In this section we describe a typical e-learning scenario in the Semantic Web context that we use for the examples in this paper. Semantic Web technologies

can improve the efficiency of e-learning in several ways. In our example we use ontologies to support finding the relevant resources. These include online courses, online books, different kinds of exercises, etc.

The Learn@WU system (<http://learn.wu-wien.ac.at/>) is an e-learning system that supports these kinds of resources and presents them on the Web. A subsystem of Learn@WU allows students to interactively try out exercises in tests which are randomly chosen from the learning resources.

Students are not aware of the learning resource ontology and are not interested how e-learning experts organize the learning resources, instead they have their own view of exercises. The interactive test system only requires this limited view of exercises, not the whole learning resource ontology.



**Fig. 1.** The learning resource ontology, the exercise ontology and possible mappings

An excerpt of the learning resource ontology is depicted on the left hand side of Figure 1. The ontology of exercise types, as depicted on the right hand side of the figure, is a typical smaller ontology that only consists of a part of the learning resources in a particular context. There is a set of equivalences between these two ontologies which have to be considered during a mapping, as depicted by dotted lines in the figure.

### 3 Creating Views in a Target Ontology

Our approach relies on the distinction between two kinds of ontologies, *source ontologies* and *target ontologies*. Source ontologies are developed by domain experts for representing the Semantic Web resources in a particular domain. In contrast, target ontologies reflect a user's perspective or a domain-specific or problem-specific view on Semantic Web resources. Note that this distinction between source and target ontologies is only conceptual; of course, any ontology can be a source ontology or target ontology as well.

Throughout this paper we use the following terms:

- Target ontology:** This is an ontology that covers a small aspect of an application domain. Examples of a target ontology from the e-learning domain are various kinds of online examples that can be automatically evaluated or are suited for a mark reader. Another example is a target ontology that contains the supported formats suitable for various delivery machines.
- Source ontology:** A (pre-existing) ontology that was developed specifically for a certain application.
- Data repository:** The metadata for the resources is defined in terms of some source ontologies. There might be multiple repositories using the same source ontology.
- View model and mapping model:** The rules to map the resources in the source ontologies to one or more target ontologies. The mapping model is specific to its source ontologies and target ontologies. Different view models can be defined on top of a mapping.
- View:** A set of resources, expressed in the target ontology. The view is created according to a mapping model and a view model.
- Query:** A user can use a view to query resources using the target ontologies solely.

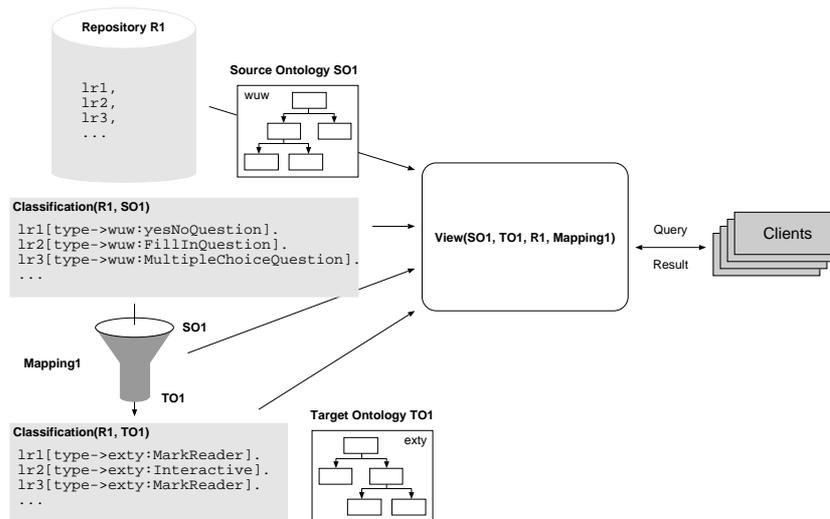
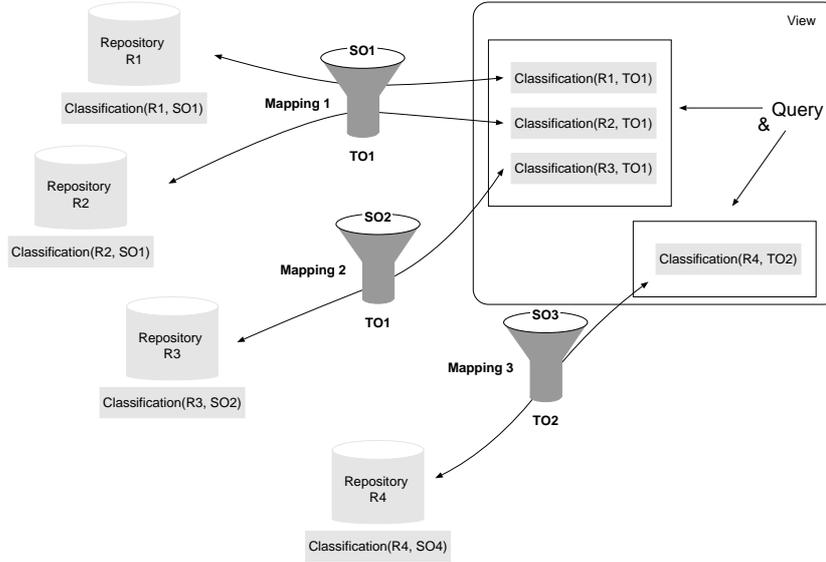


Fig. 2. Mapping Ontologies into Views

The ontology editors provide a set of *mapping rules* between the resources and the properties of the resources in the source ontologies and application ontologies. When the *view rules* are applied using the involved ontologies, mapping rules, and instance resources, a context-specific view, expressed in the application ontologies, is created. Thus clients can query the view using the application



**Fig. 3.** View with Multiple Ontologies

ontology only, without further knowledge of the more complex source ontologies. This architecture is depicted in Figure 2.

More than two ontologies can also be involved in a view. We depicted this situation in Figure 3. In the examples of Section 5 we use two ontologies only for reasons of simplicity and brevity. The examples use the language TRIPLE which is briefly explained in the next section.

## 4 Overview of TRIPLE

TRIPLE [2] is a rule language for the Semantic Web which is based on Horn logic and borrows many basic features from F-Logic [15] but is especially designed for querying and transforming RDF models.

TRIPLE can be viewed as a successor of SiLRI (Simple Logic-based RDF Interpreter [16]). One of the most important differences to F-Logic and SiLRI is that TRIPLE does not have fixed semantics for object-oriented features like classes and inheritance. Its modular architecture allows such features to be easily defined for different object-oriented and other data models like UML, Topic Maps, or RDF Schema. Description logics extensions of RDF (Schema) like OIL, DAML+OIL, and OWL that cannot be fully handled by Horn logic are provided as modules that interact with a description logic classifier, e.g. FaCT [17], resulting in a hybrid rule language.

*Namespaces and Resources* TRIPLE has special support for namespaces and resource identifiers. Namespaces are declared via clause-like constructs of the form

*nsabbrev* := *namespace*., e.g.: `rdf` := "http://www.w3.org/...rdf-syntax-ns#". Resources are written as *nsabbrev:name*, where *nsabbrev* is a namespace abbreviation and *name* is the local name of the resource. Resource abbreviations can be introduced analogously to namespace abbreviations, e.g. `isa` := `rdfs:subClassOf`.

*Statements and Molecules* Inspired by F-Logic object syntax, an RDF statement (triple) is written as: *subject*[*predicate* → *object*]. Several statements with the same subject can be abbreviated as “molecules”:

`stefan`[`hasAge` → 33; `isMarried` → yes; ...].

RDF statements (and molecules) can be nested, e.g.:

`stefan`[`marriedTo` → `birgit`[`hasAge` → 32]].

*Models* RDF models, i.e., sets of statements, are made explicit in TRIPLE (“first class citizens”).<sup>3</sup> Statements, molecules, and also Horn atoms that are true in a specific model are written as *atom@model* (similar to Flora-2 module syntax), where *atom* is a statement, molecule, or Horn atom and *model* is a model specification (i.e., a resource denoting a model), e.g.: `michael`[`hasAge` → 35]@`factsAboutDFKI`.

TRIPLE also allows Skolem functions as model specifications. Skolem functions can be used to transform one model (or several models) into a new one when used in rules (e.g., for ontology mapping/integration):

$O[P \rightarrow Q]@sf(m1, X, Y) \leftarrow \dots$

If all (or many) statements/molecules or Horn atoms in a formula (see Section 4) are from one model, the following abbreviation can be used: *formula@model*. All statements/molecules and Horn atoms in *formula* without an explicit model specification are implicitly suffixed with *@model*.

*Logical Formulae* TRIPLE uses the usual set of connectives and quantifiers for building formulae from statements/molecules and Horn atoms, i.e.,  $\wedge$ ,  $\vee$ ,  $\neg$ ,  $\forall$ ,  $\exists$ , etc.<sup>4</sup> All variables must be introduced via quantifiers, therefore marking them is not necessary (i.e., TRIPLE does not require variables to start with an uppercase letter as in Prolog).

*Clauses and Blocks* A TRIPLE clause is either a fact or a rule. Rule heads may only contain conjunctions of molecules and Horn atoms and must not contain (explicitly or implicitly) any disjunctive or negated expressions. To assert that a set of clauses is true in a specific model, a model block is used: `@model {clauses}`, or, in case the model specification is parameterized:  `$\forall Mdl$  @model(Mdl) {clauses}`.

## 5 Examples for Querying Using TRIPLE Views

In this section we provide examples for resolving typical problem scenarios in TRIPLE using the view concepts presented in Section 3.

<sup>3</sup> Note that the notion of *model* in RDF does not coincide with its use in (mathematical) logics.

<sup>4</sup> For TRIPLE programs in plain ASCII syntax, the symbols AND, OR, NOT, FORALL, EXISTS, <-, ->, etc. are used.

## 5.1 Simple Views

At first, we have to define the ontologies under our considerations. The excerpt of the Learn@WU ontology for learning resources, as introduced in Section 2 can be defined by the following TRIPLE statements:

```
@wuw:ont {
  wuw:LearningResource[rdfs:subClassOf -> rdfs:Resource].
  wuw:Book[rdfs:subClassOf -> wuw:LearningResource].
  wuw:Exercise[rdfs:subClassOf -> wuw:LearningResource].
  wuw:OpenQuestion[rdfs:subClassOf -> wuw:Exercise].
  wuw:FillInQuestion[rdfs:subClassOf -> wuw:Exercise].
  wuw:YesNoQuestion[rdfs:subClassOf -> wuw:Exercise].
  wuw:MultipleChoiceQuestion[rdfs:subClassOf -> wuw:Exercise].
  ...
}
```

In this code we define a class hierarchy of learning resources which are expressed in a TRIPLE model `wuw:ont`. We assume that the RDF Schema [11] semantics have been defined before (see appendix).

The RDF Schema rules define the `subClassOf` and `type` relationships used in the ontologies below. Using the type rules we can derive resources from the learning resource ontology:

```
question1_1[rdf:type -> wuw:OpenQuestion; wuw:difficulty -> low].
question1_2[rdf:type -> wuw:FillInQuestion;
            wuw:difficulty -> medium].
question1_3[rdf:type -> wuw:YesNoQuestion; wuw:difficulty -> high].
question2_1[rdf:type -> wuw:YesNoQuestion; wuw:difficulty -> low].
question2_2[rdf:type -> wuw:MultipleChoiceQuestion;
            wuw:difficulty -> medium].
question3_1[rdf:type -> wuw:MultipleChoiceQuestion;
            wuw:difficulty -> high].
book1[rdf:type -> wuw:Book].
...
```

Now consider there is a second ontology which defines exercise types and which is used by the web test subsystem of Lean@WU. This ontology is much simpler than the learning resource ontology and uses less resources and properties. For instance, all `difficulty` properties and all resources that are not exercises (like the book resource in the above example) can be omitted. The ontology is expressed as a second triple model:

```
@exty:ont {
  exty:LearningProgressTest[rdfs:subClassOf -> rdfs:Resource].
}
```

```

    exty:Auto[rdfs:subClassOf -> exty:LearningProgressTest].
    exty:NonAuto[rdfs:subClassOf -> exty:LearningProgressTest].
    exty:MarkReader[rdfs:subClassOf -> exty:Auto].
    exty:NonMarkReader[rdfs:subClassOf -> exty:Auto].
}

```

The concept of integrating these two ontologies is to use the `wuw` ontology as source ontology and the `exty` ontology as target ontology. Clients only use views expressed in the target ontology for the queries. We first define a mapping between the two ontologies and then define view models based on the mapping. A very simple mapping only uses the `subClassOf` relationship to model the equivalence relationships between the two ontologies. The mapping is placed in separate mapping model:

```

@exty:mappings {
    wuw:MultipleChoiceQuestion[rdfs:subClassOf -> exty:MarkReader].
    wuw:YesNoQuestion[rdfs:subClassOf -> exty:MarkReader].
    wuw:FillInQuestion[rdfs:subClassOf -> exty:Auto].
    wuw:Exercise[rdfs:subClassOf -> exty:LearningProgressTest].
}

```

For any given set of models, we can now – in a next step – define the rules for creating a view in a view model:

```

FORALL Ont1, Ont2, Mappings @view(Ont1, Ont2, Mappings) {
    FORALL R,P,O R[P -> O] <- R[P -> O]@Ont2.
    FORALL R,P,O R[P -> O] <-
        R[P -> O]@rdfschema(view(Ont1, Ont2, Mappings)).
    FORALL R,C1,C2 R[rdf:type -> C1]
        R[rdf:type -> C2]@rdfschema(Ont1) AND
        C2[rdfs:subClassOf -> C1]@Mappings.
}

```

The view model is parameterized, and takes two ontologies and a mapping between the ontologies as parameters. In this example `Ont1` is the source and `Ont2` the target ontology. The first rule states that everything in `Ont2` holds. Then, in the second rule, the RDF schema semantics, defined before, are applied to the view model. Finally we use the mapping model with its `subClassOf` relationships to create a view of the two ontologies according to the mapping.

Note that the definitions above are typically implemented by different roles, such as providers of ontologies and providers of view models. The actual resources are usually produced by the users of the system. The users (or the user tools respectively) are usually interested in creating queries as well, but here it is important that it is possible to generate queries with very simple means. For

instance, user might compose queries in GUI-based tools. The view concept helps us to limit the query syntax in such a way, and yet produce usable and user-customizable results.

For instance, a typical query using the above ontology view definition might be that the web test system requires those resources that are of type `MarkReader` (i.e. these are those exercises that it can automatically test in web forms). The system simply has to query the view for resources of the particular type and gets all resources in the view that are exercises of type `MarkReader`. Note that the simplicity of the query is due to the above view definition.

```
FORALL R <- (R[rdf:type -> exty:MarkReader])@view(wuw:ont,
  exty:ont, exty:mappings).
```

The output of the TRIPLE engine for the example above is:

```
R = question3_1
R = question2_2
R = question2_1
R = question1_3
```

## 5.2 Constraining the Views with Property Values

As already motivated in Section 2, the simple mapping rules used in the previous section might be not enough. Sensible views often have to be created under consideration of the current values of triples, for instance, property values. Consider a simple example as an extension of the scenario in the previous section: the web test system should create a shuffled mix from the set of exercises that are the basis for test questions. Further consider the test questions are defined as being either of type `FillInQuestion`, `YesNoQuestion`, or `MultipleChoiceQuestion`, and are not of the difficulty `low`. Moreover, the web test system should only use those questions from the result set that it can test automatically in web forms (i.e. those of type `MarkReader`).

We have to use a more complex inclusion rule in the mapping to implement the definition of test questions:

```
@exty:mappings {
  ...
  FORALL R R[rdf:type -> exty:TestQuestion] <-
    ((R[rdf:type -> wuw:MultipleChoiceQuestion] OR
     R[rdf:type -> wuw:YesNoQuestion] OR
     R[rdf:type -> wuw:FillInQuestion]) AND
     (R[wuw:difficulty -> medium] OR
     R[wuw:difficulty -> high]))@rdfschema(wuw:ont).
```

Note that the `exty:TestQuestion` definition is simply composed from AND and OR statements. That means it can easily be created, say, by a tool that allows users to graphically compose ontologies.

The query for the view created with this mapping again is quite simple. It simply adds the new test question type definition to the previous `MarkReader` query by composition with AND:

```
FORALL R <- (R[rdf:type -> exty:TestQuestion]
             AND R[rdf:type -> exty:MarkReader])@view(wuw:ont,
             exty:ont, exty:mappings).
```

The output of the Triple engine for this query is:

```
R = question2_2
R = question3_1
R = question1_3
```

## 6 Description Logic Specification of Mappings

An interesting class of mappings are those that can be specified with a standard description logic like *SHIQ* [9] or its Semantic Web variants (OIL, DAML+OIL, OWL).

The advantage of this approach is that users can create relatively interesting mappings with a simple point-and-click interface, since description logic expressions do not use any variables and therefore only class and property (role) names plus a small set of connectives (conjunction, disjunction, negation, ...) have to be dealt with.

The mappings are specified by connecting class expressions of the source ( $Ont_S$ ) and target ( $Ont_T$ ) ontology with the usual implication:

$$C^{Ont_S} \sqsubseteq C^{Ont_T}$$

The most simple case is where  $C^{Ont_S}$  and  $C^{Ont_T}$  are simple class names (“primitive concepts”), e.g.:

$$\text{wuw:MultipleChoiceQuestion} \sqsubseteq \text{exty:MarkReader}$$

The corresponding RDFS/OWL expression (in TRIPLE syntax) uses the `rdfs:subClassOf` to relate the two classes:

```
wuw:MultipleChoiceQuestion[rdfs:subClassOf -> exty:MarkReader].
```

The view definition `@view(Ont1, Ont2, Mappings)` contains the following rule to map instances from  $C^{Ont_S}$  to  $C^{Ont_T}$ :

```

FORALL R,C1,C2 R[rdf:type -> C2] <-
  R[rdf:type -> C1]@rdfs:schema(Ont1) AND
  C1[rdfs:subClassOf -> C2]@Mappings.

```

For the case of complex class expressions, it is much simpler to create TRIPLE rules (via an automatic mapping) instead of writing rules that handle these class expressions directly.<sup>5</sup> For expressions of the form  $C^{Ont_S} \sqsubseteq C^{Ont_T}$ ,  $C^{Ont_S}$  becomes the body and  $C^{Ont_T}$  becomes the head of a TRIPLE rule:

$$\text{TRIPLE}(C^{Ont_T}) \leftarrow \text{TRIPLE}(C^{Ont_S}).$$

As a consequence,  $C^{Ont_T}$  can only be a very simple class expression (i.e., no disjunction or negation is allowed).

As the following example shows, the mapping from the description logic expression to the corresponding TRIPLE rule is straight forward:

$$\begin{aligned}
& (\text{wuw:MultipleChoiceQuestion} \\
& \sqcup \text{wuw:YesNoQuestion} \\
& \sqcup \text{wuw:FillInQuestion}) \sqsubseteq \text{exty:TestQuestion} \\
& \sqcap \exists \text{wuw:difficulty}.\{\text{middle, high}\}
\end{aligned}$$

```

FORALL R R[rdf:type -> exty:TestQuestion] <-
  ((R[rdf:type -> wuw:MultipleChoiceQuestion] OR
   R[rdf:type -> wuw:YesNoQuestion] OR
   R[rdf:type -> wuw:FillInQuestion]) AND
  (R[wuw:difficulty -> medium] OR
   R[wuw:difficulty -> high]))@rdfs:schema(wuw:ont).

```

## 7 Related Work

Our work follows the mediator architecture suggested by Wiederhold [10]. He addresses the problem in the database context, that for end users the lack of abstraction and the need to understand the representation of data hinders the access of the available data. An other motivation for the mediator architecture was the problem of combining information from multiple databases or other information sources, where the representation and is different. In the mediator architecture a mediator plays the role of a virtual database: users might formulate the queries as if the data would be available in a mediator database and translates the queries to the data sources. The mediator synthesizes the answers from the sources and returns the answer to the user. In our approach it is possible for users to formulate their queries only in terms of the target ontology.

Several approaches have been proposed and implemented to build mediators over relational databases, for example [20].

<sup>5</sup> An alternative would be to access a description logic classifier from within the TRIPLE engine.

Calvanese et al. [8] analyze the problems of ontology integration. They describe the ontologies with Description Logics. They argue, that for capturing the mappings between ontologies the use of Description Logic is not sufficient and suggest an different notion based on queries. In our implementation the mappings build specific subset of the possible mappings we can define in description logics, but we argue that these are highly relevant in practical situations, since they are sufficient to implement a graphical tool which supports ontology mappings.

Mitra et al. [1] investigate the integration heterogeneous sources (UML, DAML+OIL). They define articulation rules that establish correspondence between concepts in different ontologies. These articulation rules are used to rewrite queries. They use a common ontology format, which is the ONION conceptual model.

The currently available ontology editors for Semantic Web also support some form of ontology mappings. Noy et al. [21] surveys ontology mapping support of these tools.

Recently many researchers investigate different aspects of ontology mappings for the Semantic Web.

Halevy et al. [14] developed a data management infrastructure for Semantic Web applications. The main focus of their work to enable interoperability of XML sources. They developed a mapping language, which is based on XQuery. The mappings are also used to support query answering, which is possible in both directions using their query answering algorithm.

On the Semantic Web the data from heterogeneous sources is described by different ontologies. Doan et al. [7] argue that in case of information processing the mapping between ontologies should happen (semi-)automatically. They developed a system that employs machine learning techniques to find the mappings.

Data integration problems and possible solutions are extensively analyzed in database literature. Wache [3] provides a classification of the problems. The schematic and data heterogeneity conflicts are described in [4]. Semantic conflicts are analyzed by Kashyap et al. [5] and Goh [6]. Structural or semantic heterogeneity or data inconsistency can cause integration conflicts:

- Structural conflicts occur if we compare the structure of the data models of different information sources. In case of structural conflicts the semantic of the data is the same, the difference is in the representation structure. Structural conflicts are for example label conflicts, where the data attributes are equivalent but different labels are used.
- In case of a semantic conflict the equivalent structural elements have to be interpreted with different semantics. A typical semantic conflict is the unit and scale conflict, when two numeric values have the same semantics, but the values have a different meaning, for example “price”, if it is represented in one information system implicitly given in Euro, while in other systems in Dollar. Another typical semantic conflict is that two information sources use different data representation types.

Due the large number of possible conflicts, often individual treatment of the problems is necessary. We only have concentrated on a subset of these problems that can be resolved with mapping rules between resources and their properties in two ontologies.

## 8 Conclusion and Future Work

On the Semantic Web, resources are described with metadata, related to some ontology. Using our approach, users or applications can express their queries for resources with respect to an ontology that reflects a user-specific or application-specific view of the domain (i.e. in the application ontologies). We demonstrated with some typical application cases from the e-learning domain how mapping rules and views can be defined using the language TRIPLE. These views allow clients to formulate queries only in the application ontologies. The mapping rules connect the instances described by the “source ontology” with the terminologies (classes and properties) of the “application ontology”. The mapping rules can be formulated in Description Logic; we demonstrated a specific class of mappings which is highly relevant in practical situations and can be expressed in TRIPLE.

Our approach has several advantages:

- Users can formulate their queries in the target ontology in a more natural way than in the source ontology. The mappings in our method are developed by ontology developers and usually not by the users who submit queries. Our motivation was not to require the user to be familiar with the domain expert’s ontology; mapping rules can be created in any case only by persons who are familiar with the source ontology.
- If the source ontology evolves, the application ontology often remains the same, only the mappings have to be updated or rewritten. In this way the users are not affected by the changes in the source ontology.
- As in other mediator approaches, our approach can be used to integrate data described in different, independently developed ontologies.
- Integration with other software systems is possible.
- Queries based on views can potentially be more efficient than queries based on the source ontologies.
- Because our implementation uses TRIPLE, the ontologies we connect with views can be represented in different ontology languages like RDFS or DAML+OIL.

Of course, our approach has some limitations:

- Some syntactic and semantic mapping conflicts (discussed in Section 7) can hardly be resolved with our approach alone.
- Our approach assumes that some mapping between different involved ontologies can be found. If it is not (easily) possible, other data integration approaches (e.g. imperative approaches) may be more appropriate. Specifically if programmatic specification of data integration is constantly required (e.g. by the end user), rule-based approaches might be cumbersome.

- The possible efficiency gain due to using a view has to be contrasted to the efficiency of creating the view. Specifically when the resources in source ontology often change, creating views can cause problems. This problem can be resolved by introducing caching structures in the view computation.

In our future research we plane to address more mapping conflicts, e.g. semantic conflicts. We also want to investigate how to integrate our approach with other data integration approaches.

## References

1. Mitra P., Wiederhold G., Decker S.: A scalable framework for the interoperation of information sources. In *Semantic Web Working Symposium*, pp. 317–329, 2001.
2. Sintek, M., Decker, S.: TRIPLE–A Query, Inference, and Transformation Language for the Semantic Web. In *Proceedings of the First International Semantic Web Conference (ISWC)*, Sardinia, June 2002.
3. Wache, H.: Semantische Mediation für heterogene Informationsquellen. Akademische Verlagsgesellschaft AKA GmbH, Berlin, Reihe “Dissertationen zur Künstlichen Intelligenz (DISKI)”, Nr. 261, 2003. In German.
4. Kim, W., Seo, J.: Classifying Schematic and Data Heterogeneity in Multidatabase Systems. *IEEE Computer*. Vol 24. No. 12. December, 1991.
5. Kashyap, V., Sheth, A. P.: Semantic and Schematic Similarities Between Database Objects: A Context-Based Approach. In *VLDB Journal* 5(4): 276-304(1996).
6. Goh, C.: Representing and Reasoning about Semantic Conflicts in Heterogeneous Information Systems. Ph.D. Thesis, MIT Sloan School of Management, 1996.
7. Doan, A., Madhavan, J., Domingos P., Halevy, A.: Learning to Map between Ontologies on the Semantic Web. In *11th International World Wide Web Conference (WWW'2002)*, Hawaii, USA, 2002.
8. Calvanese, D., De Giacomo, G., Lenzerini, M.: Ontology of Integration and Integration of Ontologies. In *Proceedings of the 2001 Description Logic Workshop (DL 2001)*.
9. Horrocks, I., Sattler, U., Tobies, S.: Practical Reasoning for Expressive Description Logics. In *Proceedings of the 6th International Conference on Logic for Programming and Automated Reasoning (LPAR'99)*, 1999.
10. Wiederhold, G.: Mediators in the Architecture of Future Information Systems. In *IEEE Computer*, Vol. 25, No. 3. March, 1992.
11. RDF Schema. <http://www.w3.org/TR/rdf-schema/>.
12. DAML+OIL. <http://www.w3.org/TR/daml+oil-reference>.
13. OWL. <http://www.w3.org/TR/owl-ref/>.
14. Halevy, A., Ives, Z., Tatarinov, I., Mork, P.: Piazza: Data Management Infrastructure for Semantic-Web Applications. In *Proceedings of the Twelfth International World Wide Web Conference*, Budapest, Hungary, May 2003.
15. Kifer, M., Lausen, G., Wu, J.: Logical Foundations of Object-Oriented and Frame-Based Languages, In *Journal of the ACM*, Jul. 1995., vol 42, pp. 741–843.
16. Decker, S., Brickley, D., Saarela J., Angele, J.: A query and inference service for RDF. In *The Query Languages Workshop, QL'98*, WorldWideWeb Consortium (W3C), Boston, USA, 1998., <http://www.w3.org/TandS/QL/QL98/>.
17. Horrocks, I.: The FaCT System. 2001., <http://www.cs.man.ac.uk/~horrocks/FaCT/>

18. Brantner, S., Enzi, T., Guth, S., Neumann, G., Simon, B.: UNIVERSAL - Design and Implementation of a Highly Flexible E-Market Place of Learning Resources, In *Proceedings of the IEEE International Conference on Advanced Learning Technologies.*, Madison (WI), USA, August, 2001.
19. Nejd, W., Wolf, B., Qu, C., Decker, S., Sintek, M.: Edutella: A P2P Networking Infrastructure Based on RDF. In *Proceedings of the 11th International World Wide Web Conference (WWW2002)*, Hawaii, USA, May 2002.
20. Chawathe, S., Garcia-Molina, H., Hammer, J., Ireland, K., Papakonstantinou, Y., Ullman, J., Widom, J.: The TSIMMIS Project: Integration of Heterogeneous Information Sources. In *Proceedings of IPSJ Conference*, pp. 7-18, Tokyo, Japan, October 1994.
21. Noy, N. F., Musen, M. A.: Evaluating Ontology-Mapping Tools: Requirements and Experience. In *Workshop on Evaluation of Ontology Tools at EKAW'02 (EON2002)*. 2002. [http://www.smi.stanford.edu/pubs/SMI\\_Reports/SMI-2002-0936.pdf](http://www.smi.stanford.edu/pubs/SMI_Reports/SMI-2002-0936.pdf).

## Complete example

```
// 0. rdf schema semantics
FORALL Mdl @rdfschema(Mdl) {
  FORALL O,P,V O[P->V] <- O[P->V]@Mdl.
  FORALL O,P,V O[P->V] <- EXISTS S
    (S[rdfs:subPropertyOf->P] AND O[S->V]).
  FORALL O,P,V O[rdfs:subClassOf->V] <-
    EXISTS W (O[rdfs:subClassOf->W] AND W[rdfs:subClassOf->V]).
  FORALL O,P,V O[rdfs:subPropertyOf->V] <-
    EXISTS W (O[rdfs:subPropertyOf->W] AND W[rdfs:subPropertyOf->V]).
  FORALL O,T O[rdf:type->T] <-
    EXISTS S (S[rdfs:subClassOf->T] AND O[rdf:type->S]).
}

// 1. learning object ontology plus some resources at WUW
// wuw := 'http://learn.wu-wien.ac.at/#'
@wuw:ont {

  // some classes for learning objects
  wuw:LearningResource[rdfs:subClassOf -> rdfs:Resource].
  wuw:Book[rdfs:subClassOf -> wuw:LearningResource].
  wuw:Exercise[rdfs:subClassOf -> wuw:LearningResource].
  wuw:OpenQuestion[rdfs:subClassOf -> wuw:Exercise].
  wuw:FillInQuestion[rdfs:subClassOf -> wuw:Exercise].
  wuw:YesNoQuestion[rdfs:subClassOf -> wuw:Exercise].
  wuw:MultipleChoiceQuestion[rdfs:subClassOf -> wuw:Exercise].

  // some learning resources
  question1_1[rdf:type -> wuw:OpenQuestion; wuw:difficulty -> low].
  question1_2[rdf:type -> wuw:FillInQuestion; wuw:difficulty -> medium].
  question1_3[rdf:type -> wuw:YesNoQuestion; wuw:difficulty -> high].
}
```

```

    question2_1[rdf:type -> wuw:YesNoQuestion; wuw:difficulty -> low].
    question2_2[rdf:type -> wuw:MultipleChoiceQuestion; wuw:difficulty -> medium].
    question3_1[rdf:type -> wuw:MultipleChoiceQuestion; wuw:difficulty -> high].
    book1[rdf:type -> wuw:Book].
}

// 2. exercise types ontology
// exty := 'http://www.exercisetype.org/#'
@exty:ont {
    exty:LearningProgressTest[rdfs:subClassOf -> rdfs:Resource].
    exty:Auto[rdfs:subClassOf -> exty:LearningProgressTest].
    exty:NonAuto[rdfs:subClassOf -> exty:LearningProgressTest].
    exty:MarkReader[rdfs:subClassOf -> exty:Auto].
    exty:NonMarkReader[rdfs:subClassOf -> exty:Auto].
}

// 3. mapping definitions
@exty:mappings {
    wuw:MultipleChoiceQuestion[rdfs:subClassOf -> exty:MarkReader].
    wuw:YesNoQuestion[rdfs:subClassOf -> exty:MarkReader].
    wuw:FillInQuestion[rdfs:subClassOf -> exty:Auto].
    wuw:Exercise[rdfs:subClassOf -> exty:LearningProgressTest].
}

// 4. view ontology
FORALL Ont1, Ont2, Mappings @view(Ont1, Ont2, Mappings) {
    FORALL R,P,O R[P -> O] <- // everything in Ont2 holds
        R[P -> O]@Ont2.

    FORALL R,P,O R[P -> O] <- // apply rdf schema semantics on "self"
        R[P -> O]@rdfschema(view(Ont1, Ont2, Mappings)).

    FORALL R,C1,C2 R[rdf:type -> C1] <- // rdfs:subClassOf mapping
        R[rdf:type -> C2]@rdfschema(Ont1) AND
        C2[rdfs:subClassOf -> C1]@Mappings.
}

// 5. query
// give me the types of all resources from the WUW ontology
// wrt. view ontology (i.e., in terms of exty ontology)
FORALL R,O <- R[rdf:type -> O]@view(wuw:ont, exty:ont, exty:mappings).

// query Mark Readers for online system
FORALL R <- (R[rdf:type -> exty:MarkReader])@view(wuw:ont, exty:ont, exty:mappings).

```