



D13.3v0.1 Mediation

WSMO Working Draft 28 June 2004

This version:

<http://www.wsmo.org/2004/d13/d13.3/v0.1/20040628>

Latest version:

<http://www.wsmo.org/2004/d13/d13.3/v0.1/>

Previous version:

<http://www.wsmo.org/2004/d13/d13.3/v0.1/20040526>

Editor:

Adrian Mocan

Authors:

Adrian Mocan
Emilia Cimpian

Reviewer:

Jos de Bruijn

This document is also available in non-normative [PDF](#) version.

Table of contents

1. Introduction

2. Problem Definition and Scope of Ontology Mediation in WSMX

2.1. Problem Definition

2.2. Addressed Ontologies

2.3. Scope of Ontology Mediation in WSMX

3. Requirement analysis

3.1. Ontology Mediation in WSMX

3.2. Ontology to Ontology Mediation Problems - Ontology Mismatches

4. Proposed Solutions

4.1. WSMX Mediation Approach

4.2. Addressed Ontology Mismatches

4.2.1. Different composition of concepts

4.2.2. Incomplete/Over-complete concept representation

4.2.3. Data-type mismatch

4.2.4. Cardinality

5. Prototype Implementation

5.1. Graphical User Interface

5.2. Mapping Rules Generator Module

5.3. Execution environment

6. Related Work

6.1. PROMPT

6.2. MAFRA

1. Introduction

This deliverable presents our approach to mediation in the Web Service Execution Environment [WSMX, 2004], an environment that is designed to allow dynamic mediation, selection and invocation of the web services. WSMX has as final scope the domain defined by Web Service Modeling Ontology (WSMO) [WSMO, 2004], in order to provide a test bed and a proof for the ideas presented in WSMO. The WSMX work includes the creation of a conceptual model [WSMX_O, 2004], the description of the execution semantics [WSMX Execution Semantics, 2004], and an architectural and software design [WSMX Architecture, 2004], together with a working implementation [WSMX Implementation, 2004]. Part of this work, this document describes the mediation problems addressed in this context, together with the appropriate solutions and a software implementation. In the first phase, the scope of WSMX is WSMO-Lite [WSMO-Lite, 2004], which offers a minimal, yet meaningful, subset of WSMO (also referred as WSMO-Standard). In the future WSMX will extend its scope to cover both WSMO-Standard and WSMO-Full [WSMO-Full, 2004]. As a consequence the mediation addressed in this document is the only form of mediation referred by WSMO-Lite - ontology to ontology mediation.

It is well known now that models or ontologies describing the same or related problem domains are created by different entities around the world. This is the reason why more and more systems and applications require mediation for being able to integrate and use heterogeneous data sources. The mapping between models are required in several classes of application, as *Information integration and Semantic web*, *Data migration* or *Ontology Merging* [Madhavan et al., 2002].

Unfortunately, it is impossible for a mapping tool to be able to automatically generate mapping rules. Some of the best results of research projects in this area are mapping tools that are able to validate or to suggest possible mappings, but in some points of the mapping process the domain expert intervention still remains a necessity. Efforts for building such tools are concentrated in two main directions [Madhavan et al., 2002]: the first is the designing of heuristics based on structure, on the non-functional properties and on names of the concepts involved, in order to generate the most plausible mappings [Noy & Musen, 2000]. Sometimes, domain independent heuristics augmented by more specific and domain related heuristics can be used [Mitra et al., 2000]. The second approach is to use machine learning techniques for semi-automatically obtaining the required mappings. For example, [Doan et al., 2002] uses the instances of one ontology for learning a classifier for it, and then applies this classifier for classifying the instances of the second ontology, and other way around. In this way, the probability for two concepts to be similar can be approximated using the fraction of instances that belong to both of them.

The mediation solution proposed by us is based on the first approach described above: we use decomposition methods for exploiting the internal structure characteristics of the concepts involved in the mediation process. In addition, strategies for suggesting the proper decompositions as well as the proper mappings between un-decomposable concepts will be provided. As for the second approach ,

the machine learning based strategies are not suited for our needs because these methods require the usage of large sets of training data (in our case instances) for obtaining a high accuracy of the predicted mappings. But the aim of the execution environment mediation is that, for a given instance (obtained at the runtime), from the source ontology, to provide the correspondent instance (or instances) from the target ontology. Additionally, one has to keep in mind that each instance that has to be mediated contains a piece of information about the internal business processes and behaviours and almost no company or enterprise will agree to grant access to its business information (i.e. all its available instances), not even for the noble scope of allowing the learning of useful mappings for its future cooperation. That is, no training data set is available, but only the instance that has to be mediated.

The entire process of mediation described in here consists of three main parts: the creation of mappings, the creation of appropriate mapping rules and the execution of the mapping rules. The first two parts are carried out at the schema level, all the mappings being created considering only the schema information, without using any instance data. Furthermore, the fact that the execution environment scope is the WSMO domain implies that the mediated ontologies are conforming to WSMO conceptual model for ontologies, having the same meta-level definition. The third part, the execution, involves also the instance data considering the given source instance and creating the target instance(s).

Although this deliverable describes the execution environment mediation, the first part of the three presented above, takes place outside the execution environment. Its implementation consists of a graphical environment that allows the expert domain to place his or her inputs and of the implementation of the required processes that assures the semi-automatic mediation behaviour of the module (e.g. the similarity computing process). The implementations of the second and the third part are included in the WSMX environment and provide means of translating the given source instances in target instances, based on the already created mappings.

This document is structured as follows: Section two presents the problem definition and scope of ontology mediation in WSMX. In this section the input ontologies are also described, introducing the conceptual model for ontologies of Web Service Modeling Ontology (WSMO). The third section presents the requirements that have to be achieved by WSMX mediation together with an exhaustive classification of potential ontology mediation mismatches. Section four introduces the terminology adopted in our approach and describes the solutions proposed for ontology mediation and the set of addressed problems. Section five briefly presents the prototype that implements the ideas presented in the previous sections. Finally, section six offers an overview of two related works in this area and section seven draws some conclusions and refers to the further directions of WSMX mediation.

2. Problem Definition and Scope of Ontology Mediation in WSMX

In this section an overview of the domain problem and of the scope of the mediation is provided. The first subsection presents the problem definition; the second subsection offers a short description of the ontologies used in our mediation process and finally the last subsection draws the scope of the ontology mediation in WSMX.

2.1. Problem definition

In the last years one of the solutions adopted for dealing with the heterogeneity over the World Wide Web was to add semantic information to the data. This approach will lead in the near future to the development of a high number of ontologies, modeling aspects of the real world. But unfortunately a new problem arose: how to integrate applications that are using different conceptualization of the same domain. Furthermore, the web services enriched with the semantics offered by the ontologies aim to enable a dynamic and a strongly decoupled communication and cooperation between different entities.

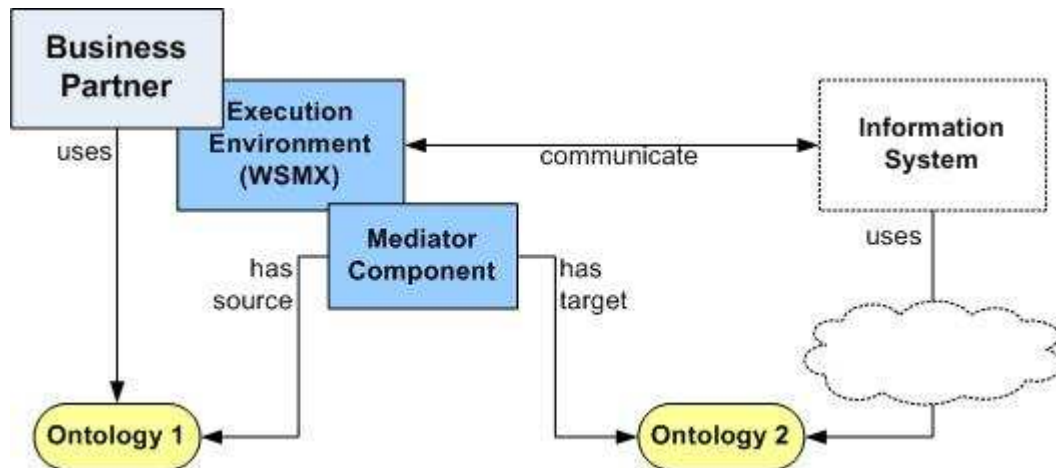


Figure 1. WSMX Mediation Scenario

As a consequence, it is very easy to imagine the following scenario (see Figure 1): two enterprises decide to become partners in business process. One of the enterprises wants to buy something that the other enterprise offers, and for achieving this a sequence of messages has to be exchanged (e.g. purchase orders and purchase order acknowledgements). Each of these messages is represented in terms of the sender's ontology, and each of the business partners (e.g. enterprises) understands only messages expressed in terms of its ontology. One of the roles of the execution environment (by mean of mediation), is to transform if necessary, the received message from terms of sender's ontology in terms of receiver's ontology, before sending it further. From the ontologies perspective each message contains instances of the source ontology that have to be transformed in instances of the target ontology.

2.2. Addressed ontologies

Web Service Modeling Ontology [WSMO, 2004] is a meta-ontology for describing several aspects related to Semantic Web Services. Its aim is to facilitate the usage of Web Services, providing means for semi-automatic discovery, invocation, composition, execution, monitoring, mediation and compensation. For achieving this ambitious goal, WSMO defines four main modeling elements: Web Services, Goals, Ontologies and Mediators. The Web Services represent the functional part - each Web Service offers a functionality which has to be semantically described, for allowing its (semi-)automatic use. A Goal specifies the objective of a client - a client consults a Web Service only if this Web Service may satisfy his goal. The Ontologies have the role of providing semantics to the information used by all the other components, and the Mediators provide interoperability among the other components.

The purpose of this deliverable is to present a possible approach for constructing a mediation system able to mediate between two ontologies (called *ooMediator* in WSMO). Our component is able to mediate between ontologies which are conforming to WSMO-Lite Specification [WSMO-Lite, 2004]. WSMO-Lite represents a minimal, yet meaningful, subset of WSMO (also referred as WSMO Standard). We choose WSMO-Lite as reference for the first version of this deliverable and for the resulting software module, but in the future we plan to extend our approach to WSMO-Standard conceptual model for ontologies. In what follows we are presenting the definitions (adopted from WSMO-Lite) of the elements that play an important role in this approach.

The first definition we have to consider is, of course, the ontology definition. As presented in *Listing 1* WSMO ontologies consist of non-functional properties, used mediators, concepts, relations and instances.

Listing 1: Ontology definition

```
ontology[
nonFunctionalProperties => nonFunctionalProperties
usedMediators =>> ooMediators
conceptDefinitions =>> conceptDefinition
relationDefinitions =>> relationDefinition
instances =>> instance
]
```

In *Listing 2* you can see that WSMO-Standard nonFunctionalProperties definition WSMO-Lite has only two elements:

Listing 2: Non-functional (core properties) definition

```
nonFunctionalProperties[
title => title
identifier => identifier
]
```

The title is actually the name by which the ontology is known and the identifier represents an unambiguous reference to the element within a given context. Every component of the WSMO ontology has the same non functional properties: *title* and *identifier*.

Regarding the usedMediators, as previously stated, the aim of this deliverable is to describe strategies and mechanisms for the construction of an *ooMediator* together with a software module that implements them.

The next element, *conceptDefinitions*, is playing a major role, the concepts and their attributes being the endpoints of the links that represents the basis of mediation process. The concept definition is presented below:

Listing 3: Concept Definition

```
conceptDefintion[
nonFunctionalProperties => nonFunctionalProperties
superConcepts =>> simpleLogicalExpression
attributes =>> attributeDefintion
]
```

As stated before, the concepts' non functional properties are the same as defined in

Listing 2. For each concept a finite number of concepts that represent its direct super concepts and a set (possible null) of attributes are defined. The attribute definition is presented in *Listing 4*:

Listing 4: Attribute Definition

```
attributeDefinition[
nonFunctionalProperties => nonFunctionalProperties
rangeDefinition => simpleLogicalExpression
]
```

The simple logical expression that stands for rangeDefinition is a type constraint that can represent in the simplest case, the extension of a concept in the ontology, or more sophisticated constraints if necessary.

The relations are not yet addressed by our mediation approach. In the future we will extend it for allowing the relation mediation too.

The instances are defined as follows:

Listing 5. Instance definition

```
instance[
nonFunctionalProperties => nonFunctionalProperties
instanceOf =>> logicalExpression
attributeValues =>> attributeValueDefinition
]
```

The instances are actually the subject of our mediation component. We define mappings between concepts of ontologies, which state for example that concept A from the first ontology correspond to concept B from the second ontology, only for being able to construct at runtime instances of B that are equivalent with given instances of A.

2.3. Scope of Ontology Mediation in WSMX

The general scope of mediation in WSMX is similar with the scope of some other components part of WSMX architecture: to provide a default implementation of the desired functionality, in our case ontology to ontology mediation. The future intension is to provide support for integration of external components, providing the same functionality, but another (or better) realization of it.

As mentioned before, the WSMX mediation addresses (for now) only ontology to ontology mediation. The intention is to provide support for transforming instances referring to one ontology in instances specified in terms of the other ontology. The transformations applied to the source ontology affect both their structure and their values - the rules that drive these transformations may be seen as a mixed of *integration* and *context* rules as described in [Wache, 1999]. The rules are build based on the similar entities from the two ontologie (i.e. concepts and attributes) identified before the runtime and saved in a storage internal to WSMX. The similar entities are defined and made known to WSMX by using a graphical user interface which allows the user to place his inputs and guides him/her through the entire process. If the similar entities are identified based on human inputs (choices, validations etc.), the rest of the process (rules generation and instances transformations) is carried out automatically, without human intervention.

3. Requirement Analysis

This section defines the main requirements for the mediation component in WSMX, and presents the main problems that need to be addressed during the ontology to ontology mediation process. In the first subsection the functionality offered by the mediation component is analyzed in more details, while in the second subsection a set of ontology mismatches is presented (as identified by [Visser et al., 1997]) relating them with our approach.

3.1. Ontology Mediation in WSMX

The ontology to ontology mediation in WSMX should provide three main functionality: the first one should facilitate the identification of the potential similarities between the entities (i.e. concepts and attributes) of two given ontology; the second one should provide the means of exploiting these similarities by creating appropriate rules to express the links between the two ontologies; the third functionality should offer an environment in which the existing links could drive the transformation process from the source to the target. In the following sections we will analyze the requirements for each of this functionalities.

Identifying similarities

This functionality can be seen as the first stage of mediation process between two ontologies. Even if the two ontologies are modeling the same aspects of the real world they could be developed by different parties, so different terminology might be used or different modeling strategies might be chosen. As a consequence we can infer the first requirement:

- the mediation process has to provide means to identify the similar entities from the two conceptualizations.

In fact, the process of finding similarities involves the identification of entities with the same semantic and, unfortunately, even if different heuristics and techniques were developed for addressing this issue, the human user inputs are still necessary. Considering this:

- the mediation component has to offer an easy way for incorporating the domain expert inputs: a graphical interface that offers a visual representation of his actions and decisions.

Further more:

- this interface should include a set of the above mentioned heuristics and techniques

for reducing the human effort from a laborious and error prone work to simple choices and validation decisions.

The graphical interface should provide facilities for:

- browsing the input ontologies;

- guiding the user through the whole process;
- hiding from the user the irrelevant information;
- offering suggestions, based on the already discovered similarities and other appropriate heuristics.

All the time the user should have access to the decisions he/she already made for potential updates and refinements. After the process ends

- the discovered similarities have to be saved in a persistent storage for further usage

either for continuing later the similarity discover between the two ontologies or for using them in the next stage of the mediation process. The way of storing these similarities should be independent from the language used for expressing the ontologies or for expressing the links in the next stage.

Creating the links

The next logical step that should be taken is to exploit the already identified similarities and to actually

- create specific links between the two ontologies.

Each link should serve a well defined purpose, which is to specify all the transformations that has to be performed for transforming a set of instances from the terms of the first ontology in the terms of the second ontology. Such a link may contain specifications of the structural transformations of the source instance and/or of the values referred by that instance - transformations dictated by the structural differences between the ontologies and the differences between the contexts in wich the instances are created [[Wache, 1999](#)].

These links may be or not expressed in the same language as the ontologies - the decision of what language to use being determined by the complexity of the transformations required and of the type of entities to be transformed. This stage of the mediation process has to be completely automatic;

- the links have to be created without any user or domain expert intervention.

Translating instances

The last step of the mediation process implies the actual transformations, and has to be performed based on the links created in the previous stage. That is, this stage (we may call it execution engine) should provide

- means for executing the transformations specified in links

and for each given source instance a set of target instances has to be created. The execution engine is directly dependent of the language chosen in the previous stage for expressing the links and should be able to

- execute the transformations in a completely automatic way without any human

intervention.

3.2. Ontology to Ontology Mediation Problems - Ontology Mismatches

In this section we present an exhaustive set of ontology mapping mismatches identified by [Visser et al., 1997] and analyze which of them should be addressed in our approach. The possible mismatches between ontologies are classified from two points of view: from the conceptualization point of view (*conceptualization mismatches*) and from the explications point of view (*explication mismatches*).

Conceptualization mismatches

The conceptualization mismatches refer to differences between the ontological concepts and between the relations among these concepts. In the first case they consist of the *categorization mismatches* and the *aggregation level mismatches*, which are described in the following paragraphs.

Categorization mismatch occurs when the two conceptualizations distinguish the same concepts but organize them in different sub-concepts. A common example could be if one conceptualization is organized around the *mammals* and *birds* concepts and the other conceptualization around *carnivores* and *herbivores*. Our approach, considering the fact that our aim is to create for the given source instance a set of target instances, is not concerned with the mediation of sub-concept relationship between concepts, but only with the identification of the most similar concept from the target, for a given source concept.

The *aggregation-level mismatches* appear because of the usage of different levels of abstraction for defining similar concepts. As an example, one ontology can use the concept of *person* and other ontology can use the concepts of *male* and *female* without having a concept similar with *person* as super-concept. As in the previous case the scope of our mediation process seems to reduce the dimension of these types of mismatches.

Another conceptualization mismatch is the *relation mismatch* and it is determined by the differences in the hierarchical relation between concepts and in the assignment of the attributes to the classes. There are three types of such mismatches identified: *structure mismatches*, *attribute-assignment mismatches* and *attribute-type mismatches*.

The *structure mismatches* are determined by the way the similar concepts are structured by means of relations. This version of mediation in WSMX doesn't address the mappings between the relations defined in the source and target ontology, even though the special case of binary relations is addressed in this approach by considering them attributes. As a consequence, these mismatches, together with the *attribute-assignment mismatches* (generated by the different assignation of attributes to concepts) and *attribute-type mismatches* (occurring when two concepts have the same assigned attribute but differ by their assumed instantiation), are covered by the WSMX mediation both in the stage of finding similarities and in the links creation process.

Explication mismatches

[Visser et al., 1997] classify these mismatches based on the assumption that an ontology is formed by a set of definitions and each definition has the following form: Def=<T, D, C> where T stands for the term used, D for the definiens and C is the ontological concept that is explained. For two given definitions (one from the source ontology and one from the target ontology) there are eight (2^3) possible combinations of mismatches, but there are two that are not interesting for this discussion: when all three terms are different (if there are no correspondences we cannot talk about mismatch) or when all three terms are identical (there is no mismatch). All six remaining mismatches should be addressed by the WSMX mediation in the similarity identifying stage, by being able to discriminate between the cases when the two analyzed definitions are referring to the same concept or not. In fact, the task of identifying similarities involves the analysis of the terms (T), and of the definiens (D), and based on this analysis a conclusion has to be drawn: the explained concepts (C) are similar or not. We present further the six possible combinations of mismatches.

- **CT mismatch** - the modeled concepts and terms used are different but the definitions are the same. In terms of our approach, this means that we have the same definitions (the same attributes) for two concepts (with different names) that model different aspects of the domain. Even if the definitions are identical the concepts modeled are different - in this case the domain expert input should be considered for taking the correct decision.
- **CD mismatch** - the modeled concepts and the definitions are different but the terms used coincide; in our context, two different concepts are modeled using different attributes. Our assumption is that all concepts from the input ontologies are different until it is specified otherwise, so the system should be able to use the appropriate heuristics for inferring the correct suggestions (based on this suggestion the domain expert should be able to make his final decisions and validations). The terms used in this case are called homonyms.
- **C mismatch** - the modeled concepts are different but the definitions and terms used are identical. We have the same problem as in the CT mismatch. The domain expert intervention in this case is a must. As a remark, for domain related ontologies this case should be very rare. As in the previous case the terms used are called homonyms.
- **TD mismatch** - the modeled concepts are similar but the terms used and the definitions are different. This means that two similar concepts are denoted by different terms and modeled using different attributes. On the first sight one may conclude that having a different definition for the two concepts the system will fail in making accurate suggestions. In reality, if the two input ontologies are modeling related domains, it is very likely that the definitions are converging to the same primitive elements (see Section 4.1 for more details related to the approach towards these requirements). The terms used in this case are called synonyms.
- **T mismatch** - the modeled concepts and their definitions are similar but the terms used are different. The offered heuristics should be able to suggest the two modeled concepts as similar (in this case the terms are also called synonyms). Based on this suggestion the domain expert should be able to make his final decisions and validations.
- **D mismatch** - the modeled concepts and the terms used are the same but their definitions are different. The situation is very similar with the TD mismatch but the identical terms used should improve the quality of the returned suggestions.

The WSMX mediation has to be able to respond to these ontology to ontology mediation problems. Some of these problems are eliminated by the nature of the

problem we intend to solve, some of them should be addressed by the mean of the heuristics provided (heuristics results may be validated by the human expert) and others only by the explicit intervention of the human user in the similarity identification process.

4. Proposed Solutions

In this section, our approach to ontology mediation, and the classes of problems we intend to address are described. In the first subsection, we introduce the terminology used in the rest of the document, and we also emphasize the solution we propose for reaching the functionalities described in the requirements section. In the second subsection, our own classification of the addressed problems is presented; this classification was chosen based on our approach's characteristics, and represents only a subset of the possible existing problems in ontology to ontology mediation, but may be seen as a starting point both from theoretical and from practical point of view. Based on the obtained results and on the experience we will gain during this work, this set of addressed problems will be extended in order to cover a larger set of possible situations.

4.1. WSMX Mediation Approach

[Wiederhold, 1994] identifies three main approaches that deal with ontology integration. The first of these approaches implies the use of domain experts for creating definitions for the terms in the used ontologies, definitions that must be accepted by all parties involved. When these definitions are completely documented and released, all the parties have to conform and adjust to these definitions. The second approach is to assume that the terms from the subject ontologies are matching and further discovered mismatches are resolved by annotating the corresponding terms with some source or domain identifiers. The last approach, on a contrary, assumes that all the terms from the given ontologies are distinct and are denoting different concepts while it is not specified otherwise. This kind of information is explicitly stated using mapping rules.

The first approach is used in [Dou et al., 2003] in the design of the OntoMerge system used for ontology merging and automated reasoning. The system uses a merged ontology (containing symbols and facts from the source and target ontologies, together with a set of bridging axioms) for performing translations between two ontologies. But this approach is not suited for our needs, our aim being to create a tool that provides the mapping rules in a semiautomatic manner and then executes them.

The second approach could be useful when knowing that the subject ontologies are using very similar terminology for modeling the problem domain. Unfortunately the ontologies are usually created independently, so there could be major differences even in representing the same domain. So, the best choice for our mediation component remains the third one: all the concepts from the source and the target ontology are different until mapping rules specify otherwise.

Following the requirements, our solution proposes three main steps for the WSMX mediation: *mapping generation* step for dealing with similarity identification, *mapping rules creation* step for creating the links and the *execution* step for performing the instance translation. Each of these steps is presented in details in the following sections.

Mapping Generation

The first step is the most important one as well as the most challenging one. This step is concerned with the identifications of the possible links between the elements from the source and target ontologies, considering only the schema level of the input ontologies. These links are called *mappings* and each of these mappings connects either a pair of concepts or a pair of attributes. The semantic of such a concepts pair (*mapped concepts*) is that the involved concepts are modeling similar or related concepts in the source and target ontology, respectively. Each attributes pair (*mapped attributes*) denotes a way of realizing the mapping between the concepts that provides them (see section 2.2. or [WSMO-Lite, 2004] for more details about concept and attribute definitions). Even if it was the subject of intense research activities in the last ten years, the process of determining the possible mappings between two ontologies could not be totally automated. That is, the human user input is still required at least in some of the key points of the mapping process. The challenge is how to minimize the user's (usually a domain expert) effort and to reduce it to simple choices and validations.

The approach adopted in this direction aims to make full use of the advantages offered by the assumption that all input ontologies are conforming to the WSMO-Lite conceptual model for ontologies (see section 2.2). This assumption allows the defining of strategies and mechanisms applicable to all concepts and to all attributes from both the source and the target ontologies.

Most of the strategies used in this step are based on the so called *decomposition* mechanism. But for being able to properly describe this mechanism it is necessarily to introduce several others notions. We define the *source context* (*target context*) a set of elements (only attributes or only concepts) from the source (target) ontology that could be chosen to be part of a mapping in a certain point of the mapping creation process. That is, in each stage a source and a target context are defined and all the new mappings can be created only using a pair from these contexts. Further, we call a *compound concept* a concept that refers other concepts through its attributes or relations. In this stage of our research, we consider only the attributes, so compound concepts are considered to be those concepts that have at least one attribute defined. The remaining concepts are called *primitive concepts* - concepts that are no further defined and don't contain any attributes (e.g. data types like string or integer). Accordingly, a *compound attribute* is an attribute that refers through its range a compound concept and a *primitive attribute* is an attribute that refers in the same way a primitive concept. As a consequence, we have *direct mappings* when both participants are primitive concepts or primitive attributes and *indirect mappings* otherwise. In fact, the last type of mappings is determined by zero or more other indirect mappings and one or more direct mappings (we'll see further in which way). All the mappings are considered in this document to be one way mappings, even if in some particular cases they are bi-directional.

The decomposition mechanism (implemented by a recursive process) is applied each time a potential mapping has to be done. In each recursion step the contexts are updated accordingly and the decomposition process is further applied for other pair of elements chosen from the new context (the way in which the new pair is chosen is discussed later in this section). The recursive calls end when a direct mapping, or a mapping between two attributes with already mapped concept ranges, is done or when no other mappings are possible (or desired) in a certain step. When returning from recursion after at least one successful mapping an indirect mapping between the chosen pair of elements is registered. In Table 1 the possible

combinations of elements to be mapped are presented.

Table 1. Possible Mapping Situations

Source element	Target Element	Action
Primitive Concept	Primitive Concept	A direct mapping can be done and the decomposition process returns to the upper level (to the previous contexts) if there is one.
Primitive Concept	Compound Concept	This is an invalid combination because the source and the target concepts are considered to be part of different mapping levels. This combination would introduce artificial mappings (with no semantics) between primitive concepts and all the compound concepts that referes by mean of their attributes to these primitive concetps. For example, <i>String</i> could be mapped with any compound concept A that has an attribute with the range <i>String</i> and also, in a recursive manner, with any other compound concept that has at least one attribute with the range mappable with <i>String</i> .
Compound Concept	Primitive Concept	This is an invalid combination (identical with the previously described one) because the source and the target concepts are considered to be part of different mapping levels. This two combinations are covered by the <i>Compound Concept to Compound Concept</i> mappings, our <i>Primitive Concept</i> being the range of an attribute owned by one of the compound concepts.
Compound Concept	Compound Concept	The decomposition must be applied both on the source and on the target element, as a consequence, the contexts being updated in the following way: the source (target) context will be formed from the attributes owned by the source (target) concept. The decomposition process is applied further on a new pair of elements from the new source and target contexts respectively.
Primitive Attribute	Primitive Attribute	A direct mapping between these two elements is possible. In fact, in this case two steps are performed: the source and target context are updated to contain the range concept of the source attribute and the range concept of the target attribute respectively, and then primitive concept to primitive concept mapping is required. Then the decomposition process returns to the upper level (to the previous contexts) if there is one.
Primitive Attribute	Compound Attribute	The decomposition has to be applied on the target element. In this case the source context is updated to contain only the source attribute and the target context to contain all the attributes of the range concept of the target attribute. The decomposition process is applied further on a new pair of elements from the new source and target contexts respectively.
		This situation is similar with the one presented above,

Compound Attribute	Primitive Attribute	only the roles of the source and target elements are exchanged.
Compound Attribute	Compound Attribute	The decomposition process is applied to both source and target elements. The new contexts will contain the attributes owned by the ranges of the two attributes respectively. The decomposition process is applied further on a new pair of elements from the new source and target contexts respectively.

The pair of elements that is further considered in each decomposition step can be chosen in different ways. The simplest one is to ask for the domain expert inputs and in this case the decomposition process has the role of offering guidance through the mapping definition process. Another possibility is to use the same decomposition process for computing suggestions based on the internal structure of decomposable elements and on lexical similarities of the names and non-functional properties. That is, the decomposition can be applied for any two concepts from the source and target contexts respectively, and based on the number of direct and indirect mappings together with the lexical similarities, a global similarity measure can be computed. The pair with the higher similarity measure will be picked for the next step of the original decomposition process. One has to keep in mind that the computed suggestion could represent in some cases not exactly what was desired, so the human expert intervention is still necessary (at least for validations). Even if the last solution is used, there are some cases when no valid suggestions are made (a valid suggestion is one that offers a similarity degree greater than a given threshold), the human user intervention still remaining a must. For this second approach, the decomposition process together with the computed similarities measures aim to reduce the human effort, transforming a laborious and error prone process to simple choices and validations. The general decomposition algorithm is outlined in Listing 6.

Listing 6. Decomposition Algorithm

```

decomposeConcepts(Concept A, Concept B){
    if (primitive(A) and primitive(B) then
        return;

    if ((primitive(A) and not(primitive(B)) or (not(primitive
(A)) and primitive(B)) then
        throw IllegalCombination(A,B);

    ContextA = getContext(A);
    ContextB = getContext(B);
    PairsCollection = choosePairsForDecomposition
(ContextA,ContextB);
    for each P in PairsCollection do
        decomposeAttributes(P.getSourceAttribute,
P.getTargetAttribute);
}

decomposeAttributes(Attribute A, Attribute B){
    if (primitive(A) and primitive(B) then
        decomposeConcepts(A.getRange; B.getRange);

```

```

    if (primitive(A) and not(primitive(B)) then{
        PairsCollection = choosePairsForDecomposition
({A},B.getRange.getContext);
        for each P in PairsCollection do
            decomposeAttributes(P.getSourceAttribute.
P.getTargetAttribute);
        return;
    }

    if (not(primitive(A)) and primitive(B)) then{
        PairsCollection = choosePairsForDecomposition
(A.getRange.getContext,{B});
        for each P in PairsCollection do
            decomposeAttributes(P.getSourceAttribute.
P.getTargetAttribute);
        return;
    }

    PairsCollection = choosePairsForDecomposition
(A.getContext,B.getContext);
    for each P in PairsCollection do
        decomposeConcepts(P.getSourceConcepts.
P.getTargetConcepts);
}

```

One can notice that in the case of a recursive definition of a concept (e.g. contains an attribute that has as range the concept itself) this algorithm will not terminate. A simple solution that can be adopted is to consider the attribute that introduces recurrence only a finite number of times in the decomposition process. This solution is also adopted in the system that generates the source instances to be mediated; the problem remains to determine this number and exactly which of the attributes of the initial concept to consider for the last recursive step, in order to match the source instance. For now, the solution adopted is to allow only one recursive step, eliminating from the second step the attribute that introduces the recursion.

All the mappings, direct or indirect, are saved in an external storage in order to be retrieved later, either for the suggestion mapping process or for the mapping rules creation.

Mapping rules creation

After the schema level mappings between the two ontologies are done, these mappings are used for the creation of the *mapping rules*. A mapping rule specifies the way in which the values in the first instance are extracted and used for the creation of the target instance(s). When there is a demand for converting a source instance in the corresponding target instance(s), the owner of the source instance is identified and then a search in the storage is performed to determine if there are some available mappings. If no mappings are available the mapping rule can not be created and in this case a return to the previous step is required. If for the given source concept one or more mapped concepts are found the similarity measures are used again for determining the most probable mapping. And finally, if there is an indirect mapping between the chosen pair of concepts, a simplified decomposition process is applied for determining all the indirect and direct mappings that lead to

this mapping. In fact, all the involved indirect mappings are used for specifying the structure of the target instance(s) that has (have) to be created and the direct mappings are dealing with specifying the passing of the actual data to the target instance(s).

The language in which the mapping rules are generated is usually determined by the language in which the input ontologies and source instance are expressed; also the new target instance will be expressed in the same language. Some translators may be used for allowing the usage of different languages (e.g. for expressing the target instance in another language).

Once a mapping rule is created, it can be saved for other potential usages or it can be redone each time this is necessary. If the used ontologies are unlikely to change, then may be much efficient to store the mapping rules in a persistent storage and just reused them when they are needed. If the ontologies are evolving and they are frequently changing, the second approach would be more appropriate - the consistency have to be maintained only between the mappings in order to have consistent mapping rules.

Execution

This step deals with the execution of the existing mapping rules using an environment that understands the language in which the source instance and the rules are represented. The result of the execution could be one or more instances of concept(s) from the target ontology. This step may also include the checking of the potential constrains defined in the source or in the target ontology and which affect the mapped elements. That is, using the execution environment one can check the correctness and the consistency of the mapping rules.

4.2. Addressed Ontology Mismatches

In this section the set of problems that are addressed in our approach are presented. In the future, this set of problems will be extended and the next versions of this deliverable and of the provided prototype will offer solutions for a greater number of possible mismatches that can appear during the mapping process.

4.2.1. Different Composition of Concepts

This is the most common problem that can appear during the mapping process and it is caused by the fact that concepts that share the same semantic meaning are modeled in different ways, e.g. using a different number and types of attributes and relations. That is, we can have different levels of conceptualizations used for modeling the same aspect of the domain, even if in the end, at the lowest level, we have the same primitive concepts. This problem is also described as granularity of match in or simply as granularity in [[Chalupsky, 2000](#)].

This problem is addressed by applying the decomposition process, revealing the internal structure of the concepts, and recursively moving from one level to another till the primitive elements are encountered. What remains to be done is to create the direct mappings between these primitive elements - we will see later that this could be in some case a non-trivial task at all.

4.2.2. Incomplete/Over-complete Concept Representation

Another problem that may appear during the mapping process is that a required concept (either from the source or from the target ontology), does not have a similar concept in the other ontology. In this case a decision has to be made regarding the missing concept and this decision mostly depends whether the missing concepts appears in the source or in the target ontology. In the first case the data will be mediated without any content loss (even if from the target point of view the data will be incomplete), but in the second case some of the content will be lost in the mediation process (but from the target point of view the data will be complete).

In general, for this kind of problems several solutions could be adopted, depending on the requirements of the two business partners. When the missing concept is on the source side the mapping rule must assure that in the created instance (target instance) a default value is filled. Also there are some cases when simply ignoring the missing concept could be enough, the responsibility of dealing with the missing value being delegated to the application that is using the new created instances. When the missing concept is on the target side the problem is more complicated. Normally, the target doesn't know what to do with an instance of a concept that it cannot understand, but there may be some situations in which the target will have to return back that instance, for example, as a participation proof in the previous communication. So, a solution could be to simply discard, during the mapping rule execution, all the instances of an unmapped concept (assuming that it will be no need for these instances in the future) or mapping them as instances of a bag concept whose only role is to gather all unmapped instances.

The simplest solution is adopted for the first version of this deliverable: when the missing concept is in the source ontology no instance of the target concept is created; when the missing concept is on the target the corresponding source instance (or value) is discarded.

4.2.3. Data-Type Mismatch

Another common problem for the mediation process appears when concepts denoting the same aspects are defined using different data types. The most common example is when the data-type for an attribute is *string* in one of the ontologies, and for the similar one in the other ontology a most specific data type is used, like *integer*.

We could have different dimensions of this problem, depending of the casting compatibility of the two data types involved. In the above example the solution is very simple when the attribute having as data type an *integer* is part of the source ontology (*integer* can always be converted to a *string*) but if we consider the other way around, the solution is not trivial at all, since not any *string* can be converted to an *integer*.

Considering our approach, this problem can appear only in the final phase of the mediation process, when mapping between primitive concepts, after all the required decompositions were already done. Between the attributes, the mappings can be simple assignments from the source attributes to the target attributes, but also more complicated formulas that make the explicit conversion from one data format to the other. The checking if the mapping rule has specified the correct operation for conversing from one data format to the other is actually done only when the rule is executed. The operations performing the required transformations can be seen as the context rules, described in [Wache, 1999].

4.2.4. Cardinality

When creating the mappings between ontologies, the following situation may also occur: some concepts from one of the ontologies are subsumed by one concept from the other ontology. In [Omelayenko & Fensel, 2001] this class of problems is referred as being 1:n, n:1 and n:m mapping cases and in [Rahm & Bernstein, 2001] it is considered to be part of the match cardinality process as set-oriented mapping case.

It is important to mention here that this problem appears only when mapping primitive elements, so no other concept (and implicitly instance) details are available. The cardinality problem complexity is dependent of where the subsumed concepts are placed: in the source ontology or in the target ontology. In the first case one can get from the subsumed concepts to the target concept by applying a certain operation (e.g. concatenation). In the second case it is not so trivial to create a mapping rule that splits a given instance (which does not have an explicitly defined structure) in a set of subsumed concepts instances. The only solution for this situation is to find out details about source instance's internal structure from other information sources (for example the instance provider). The most general case of subsumption is the n:m mapping, but it can be seen as m cases of 1:n mappings.

The problem presented in the above section is a particular case of the cardinality problem, from the adopted solutions' point of view. That is, for this case, we have to provide operations that are able to accept as input one or more source instances and to return one or more target instances. But one has to consider the fact that each of the owners might create their instances in their own way and it is the role of the mediation process to make the required transformations.

5. Prototype Implementation

As described above, the mediation process we propose consists of three main steps: *mapping generation*, *mapping rules creation* and *execution*. Accordingly, the mediation module will consist of three main sub-modules or components: a graphical user interface for defining the mappings, a component that reads the mappings from the storage and generates the appropriate mapping rules and an environment that provides the means for executing the mapping rules. Figure 2 shows how these components interact.

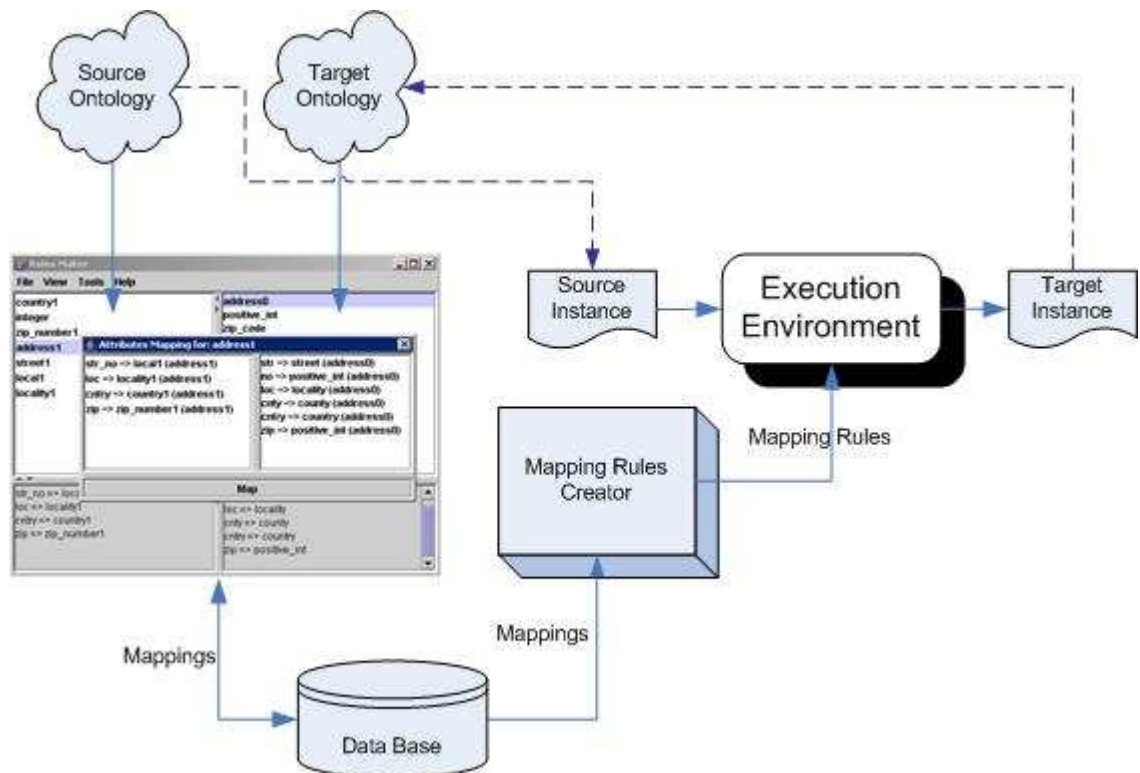


Figure 2. Overview of the Data Mediation Module

5.1. Graphical User Interface

The mapping creation process cannot be performed 100% automatically and as a consequence, the domain expert has to provide his inputs in various stages of this process. Of course, the mappings could be done entirely manually, but this proved to be a very laborious, error prone and time consuming process. Considering this, the mapping tools offer graphical interfaces in order to help and to assist the user in the mapping creation and to reduce his effort to simple choices and validations. This kind of graphical mapping tools are also described in [Maedche et al., 2002] and [Noy & Musen, 2000].

The graphical component described in here has two main roles: firstly to guide the user to the mapping creation process, and secondly, to compute the suggestions based on already done direct and/or indirect mappings. The first task is accomplished by presenting the user the current source and target contexts containing the eligible elements for mapping. In fact, in this top-down approach, the user knows what elements he wants to map and through the decomposition process the contexts are updated in each step guiding the user to the primitive elements that will be mapped in the last step. This top-down approach may be very easily combined with a bottom up approach (this being the recommended solutions), in order to enable the suggestion mechanism to return better results. That is, the user can address in the beginning the mapping of some of the basic elements (including primitive elements), in this way constructing a basic vocabulary which will be referred in the next top-down mappings.

The second task is achieved by computing a similarity factor between the elements to be mapped. This similarity measure is based on two similarity factors: the first one is a heuristic function that computes an eligibility factor based on the already done mappings (direct or indirect mappings), and the second one is a lexical similarity

factor based on the syntactical similitude between the name of the elements and their nonfunctional properties. These two factors are used together with the decomposition process in order to compute the similarities between compound elements as well as between primitive elements. It is important to notice that without establishing a basic set of mappings between primitive and basic elements of the two ontologies, the results returned by the first factor, the eligibility measure, cannot be accurate. For this, the mappings between primitive concepts could be determined by using the lexical factor, together with a thesaurus (e.g. WordNet).

The graphical interface is built in Java and is able to load Flora-2 [Flora-2] ontologies. The created mappings are stored in an external storage, in this case a relational data base, from where they can be loaded by mapping rules generator module. Also, by means of this graphical interface the user can load already existing mappings from the external storage for further refinements or as support in computing the suggestions.

5.2. Mapping Rules Generator Module

This module deals with the creation of the mapping rules and expresses them in a format that can be later executed. It takes as input a concept from the source ontology and an identifier of the target ontology and search the storage for finding mappings between the given concept and concepts from the target ontology. There are three situations that can occur: no mappings to target concepts are found, exactly one mapping to a concept from the target is found, or several mappings to different concepts from the target ontology are identified.

In the first case the creation of rule fails and an error message is returned. This means that the return to the previous step is required in order to provide the suited mappings. The second case, when only one mapping exists, is the easiest to solve one: the decomposition process is applied in order to consider all involved elements (attributes and concepts). But the third case is a little bit more complicated. Having more than one mapping could mean that there are several concepts in the target that are similar with the concept in the source ontology. Going further, and assuming that the mappings are correct, this means that in the target, some equivalent concepts are defined which usually is not a good modeling choice, unless this is explicitly specified (e.g. by an inheritance relationships, or by subsumptions and aggregations). We will consider that the input ontologies are correctly modeled, so the source concept can be mapped to more equivalent concepts with a different degree of generality. As a consequence, a decomposition process is applied and a similarity factor is computed based on the indirect and direct mappings implied by the initial ones. As a remark, even if the target ontology is badly modeled and there are equivalent concepts without an explicit relation between them, the above described process will pick one, and due to these equivalences the choice will be as good as any other.

The generated mapping rules will be taken by the execution environment and optionally they could be stored for performance improvements or caching facilities. In our case the mapping rules are represented in Flora2 and no storing actions are performed.

5.3. Execution environment

As we mentioned above the mapping rules are expressed in Flora so the execution

environment is a Flora2 environment for XSB [Flora2]. This environment is wrapped by a Java interface and it is ready to take the source instance and the mapping rule and to return the target instance(s) by executing the rule. The target instance is also expressed in Flora but an XML description for it may be provided. For this version no other computations are done in this environment but it could be also used for testing the validity of the mappings and perform constraints checking.

The last two sub-modules are used by WSMX at runtime to provide and to execute the mapping rules. For each source instance provided by WSMX, the available mappings are retrieved from the storage and a mapping rule is created. The new created mapping rule is forwarded to the execution module, executed, and the result (i.e. the new target instance) is returned to WSMX. The graphical user interface is previously used, before runtime, for defining and storing the mappings between similar entities.

The first version of this prototype allows the user (more or less a domain expert) to create a set of mappings between concepts and attributes of two ontologies. The types of mismatches that can be addressed and solved using this prototype were described in section 4.2.1, 4.2.2, 4.2.3, 4.2.4. The process of creating mappings could be seen as an iterative one: the user should analyze the offered suggestion, implement some of them and then re-explore the new set of suggestion offered (the suggestions are strongly affected by the created mappings: better suggestions are obtained based on a higher number of mappings). Once the mapping process ended the mappings can be dropped in the external storage for being picked by the runtime mediation component, transformed in mapping rules and executed in order to obtain the target instances.

6. Related Work

In this section we present two related approach to ontology to ontology mediation. We chose to present PROMPT [Noy & Musen, 2000] and MAFRA [Maedche et al., 2002], because of the existing similarities with our approach both in the conceptual framework proposed and in the software implementation they offer.

6.1. PROMPT

PROMPT is an algorithm and a tool proposed by Noy and Musen [Noy & Musen, 2000] which allows semi-automated ontology merging and alignment. It takes as inputs two ontologies and guides the user through an iterative process for obtaining a merged ontology as an output. This process starts with the identification of the classes with similar names and provides a list with initial matches. Then the following steps are repeated several times: the user selects an action (by choosing a suggestion or by editing directly the merged ontology) and the tool computes new suggestion and determines the eventual conflicts. This iterative process is presented in Figure 3.

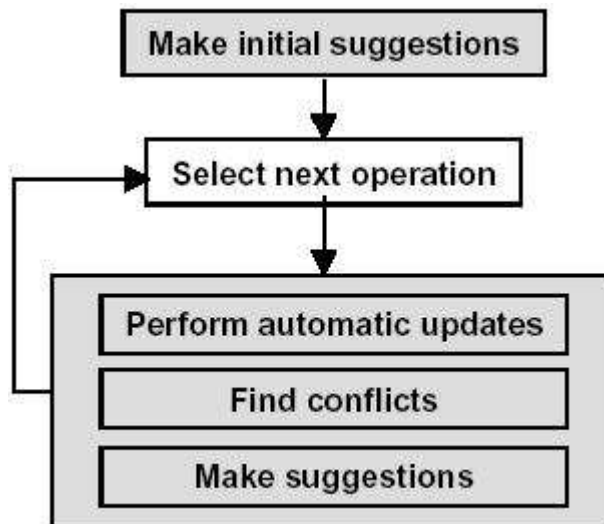


Figure 3. The PROMPT Algorithm

PROMPT is able to handle ontologies conforming to OKBC [Chaudri et al., 1998] format having at the top level classes, slots, facets, and instances. The operation allowed include merge classes, merge slots, merge binding between a slot and a class, perform deep/shallow copy of a class from one ontology to another (including/not including all the parents of the class up to the root of hierarchy and all the classes and slots it refers to). Some of the conflicts identified by PROMPT are name conflicts, dangling references (a frame refers to another frame that doesn't exist), redundancy in the class hierarchy (more than one path from a class to a parent other than root) or slot-value restriction that violates class inheritance.

Implementation

PROMPT was implemented as a plugin in Protege-2000 - a knowledge base acquisition tool, and by this PROMPT can take advantage of all the ontology engineering capabilities of this tool. Some of the feature offered by this implementation are: setting the preferred ontology (for automatic conflicts solving in favor of one ontology), maintaining the user's focus (suggestions the user sees first should relate with the frame in the same area), providing explanations for the suggestions made, logging and reapplying operations.

6.2. MAFRA

MAFRA [Maedche et al., 2002] is a Mapping Framework for Distributed Ontologies, designed to offer support in all parts of the ontology mapping life-cycle. Also a partial implementation is offered for this framework as part of the Kaon Ontology and Semantic Web Framework.

The architecture of MAFRA conceptual framework is presented in Figure 4.

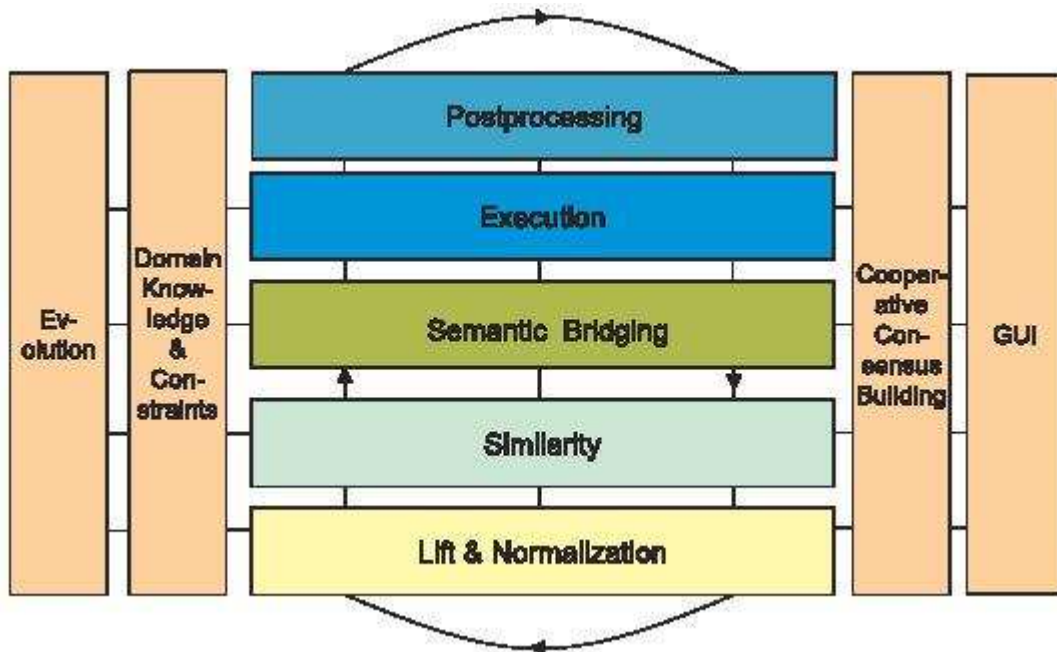


Figure 4. Conceptual Architecture

The framework is organized in two dimensions: it contains horizontal and vertical modules. The horizontal modules (*Lift & Normalization*, *Similarity*, *Semantic Bridging*, *Execution*, *Post-processing*) describe fundamental and distinct phases in the mapping process, while the vertical modules (*Evolution*, *Domain Constraints & Background Knowledge*, *Cooperative Consensus Building*, *Graphical User Interface*) run along the entire mapping process interacting with the horizontal modules.

Horizontal Dimension

The *Lift & Normalization* module has the role of coping with syntactical, structural and language heterogeneity, raising all the data to be mapped to the same representation language. Both ontologies are normalized to a uniform representation, in this case RDF(S). *Similarity* module implements strategies for discovering the similarities between entities. It determines lexical similarity, as well as so called property similarities - similarities between concepts based on their properties, either attributes or relations. Also it defines bottom-up/top-down similarities for propagating the similarities or dissimilarities from lower/upper parts of the taxonomy to the upper/lower concepts. *Semantic bridging*, based on the similarities previously determined, establishes correspondences between entities from the source and target ontology. It specifies bridges between entities in a way that each instance represented according to the source ontology is translated in to the most similar instance described according to the target ontology. The *Execution* module actually transforms instances from the source ontology to the target ontology, by evaluating the semantic bridges defined in previous step, and the *Post-processing* module takes the results for checking and improving the quality of the transformations.

Vertical Dimension

Evolution focuses on keeping the bridges obtained by the Semantic Bridging module up to date, consistent with the transformation and the evolutions that may take place in the source and target ontology. *Cooperative Consensus Building* is responsible for establishing a consensus between two communities relative to the semantic

bridges used. The need for this module comes from the multiple bridges available for performing transformations and from the aim of reducing the human contribution in the mapping process. *Domain Constraints & Background Knowledge* may improve substantially the quality of bridges by using background knowledge and domain constraints, as glossaries, lexical ontologies or thesauri. Finally, *Graphical User Interface* is required due to the fact that the mapping creation is a difficult and a time consuming process, thus extensive graphical support must be provided.

7. Conclusions and Further Directions

This deliverable provides solutions for ontology to ontology mediation problems, as starting point for the more complex problem of mediation in the context of Semantic Web Services. Our intentions are to provide conceptual foundations for a semi-automatic approach to ontology to ontology mediation together with a software implementation as a testbed and a proof of their validity. Also, the software implementation is integrated in Web Service Execution Environment (WSMX) as a default implementation of the mediation component.

As future work, our focus will be to improve the solutions proposed for ontology mediation including the extension of the problems set we address together with the creation of better heuristics for computing suggestions. Another aim is to further explore the mediation problems in context of WSMO, offering solutions for the other types of mediators which were not address in this document and integrating their implementation in WSMX mediation component.

References

[Chalupsky, 2000] H. Chalupsky: OntoMorph: a translation system for symbolic knowledge. In *Proceedings of 7th international conference on knowledge representation and reasoning (KR), Breckenridge, (CO US), pages 471-482, 2000.*

[Chaudhri et al., 1998] V. K. Chaudhri, A. Farquhar, R. Fikes, P. D. Karp, J. P. Rice: OKBC: A Programmatic Foundation for Knowledge Base Interoperability. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)* (pp. 600-607). Madison, Wisconsin, USA: MIT Press, 1998.

[Doan et al., 2002] A. Doan, J. Madhavan, P. Domingos, A. Halevy: Learning to map between Ontologies on the Semantic Web. In *WWW2002*, 2002.

[Dou et al., 2003] D. Dou, D. McDermott, P. Qi: Ontology Translation on the Semantic Web. In *ODBASE'03*, 2003.

[Fensel & Bussler, 2002] D. Fensel, C. Bussler: The web service modeling framework wsmf. In *White Paper and Internal Report Vrije Universiteit Amsterdam*, 2002.

[Flora-2] Available at <http://flora.sourceforge.net/>

[Lin et al., 2001] H. Lin, T. Risch T. Katchaounov: Adaptive data mediation over XML data. In *Journal of Applied System Studies (JASS), Cambridge International Science Publishing*, 2001.

[Madhavan et al., 2002] J. Madhavan, P. A. Bernstein, P. Domingos, A. Y. Halevy: Representing and reasoning about mappings between domain models. *Eighteenth national conference on Artificial intelligence*, p.80-86, Edmonton, Alberta, Canada, July 28-August 01, 2002.

[Maedche et al., 2002] A. Maedche, B. Motik, N. Silva, R. Volz: MAFRA - A Mapping Framework for Distributed Ontologies. In *Proceedings of the 13th European Conference on Knowledge Engineering and Knowledge Management EKAW*, Madrid, Spain, September 2002.

[Mitra et al., 2000] P. Mitra, G. Wiederhold, M. Kersten: A Graph-Oriented Model for Articulation of Ontology Interdependencies. In *Proceeding Extending DataBase Technologies, Lecture Notes in Computer Science*, vol. 1777, pp. 86-100, Springer, Berlin Heidelberg New York, 2000.

[Noy & Musen, 2000] N. Noy, M. Musen. PROMPT: Algorithm and Tool for Automated Ontology Merging and Alignment. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2000.

[Omelayenko & Fensel, 2001] B. Omelayenko, D. Fensel: An Analysis of Integration Problems of XML-Based Catalogues for B2B E-commerce. In *Proceedings of the 9th IFIP 2.6 Working Conference on Database (DS-9)*, Semantic Issues in e-commerce Systems, Hong Kong, April 2001.

[Rahm & Bernstein, 2001] E. Rahm, P. Bernstein: A survey of approaches to automatic schema matching. *The VLDB Journal* 10: 334-350, 2001.

[Visser et al., 1997] P. R. S. Visser, D. M. Jones, T. J. M. Bench-Capon, M. J. R. Shave: An analysis of ontological mismatches: Heterogeneity versus interoperability. In *AAAI 1997 Spring Symposium on Ontological Engineering*, Stanford, USA, 1997.

[Wache, 1999] H. Wache: Towards rule-based context transformation in mediators. In *S. Conrad, W. Hasselbring, and G. Saake, editors, International Workshop on Engineering Federated Information Systems (EFIS 99)*, Kuhlungsborn, Germany, 1999.

[Wache et al., 2001] H. Wache, T. Vogele, U. Visser, H. Stuckenschmidt, G. Schuster, H. Neumann, S. Hubner: Ontology-Based Integration of Information - A Survey of Existing Approaches. In *Proceedings of IJCAI-01 Workshop: Ontologies and Information sharing*, Vol. pp. 108-117, Seattle, WA, 2001.

[Wiederhold, 1994] G. Wiederhold: Interoperation, Mediation, and Ontologies, In *Proceedings International Symposium on Fifth Generation Computer Systems (FGCS94), Workshop on heterogeneous Cooperative Knowledge-Bases*, Vol. W3, pages 33-48, Tokyo, Japan, Dec. 1994.

[WSMO-Lite, 2004] D. Roman, U. Keller, H. Lausen (eds.): Web Service Modeling Ontology - Lite (WSMO-Lite), version 0.2, 2004, available at <http://www.wsmo.org/2004/d11/v0.2/>

[WSMO-Full, 2004] C. Preist, D. Roman (eds.): Web Service Modeling Ontology - Full (WSMO-Full), WSMO Working Draft v01, 2004, available at <http://www.wsmo.org/2004/d12/v0.1/20040624/>

[WSMO, 2004] D. Roman, U. Keller, H. Lausen (eds.): Web Service Modeling Ontology - Standard (WSMO - Standard), version 0.2, 2004, available at <http://www.wsmo.org/2004/d2/v0.2/>

[WSMX, 2004] E. Oren, M. Zaremba, M. Moran: Overview and Scope of WSMX, WSMO Working Draft v01, 2004, available at <http://www.wsmo.org/2004/d13/d13.0/v0.1/20040611/>

[WSMX_O, 2004] E. Cimpian, A. Mocan, M. Moran, E. Oren, M. Zaremba: Web Service Modeling Execution Environment – Conceptual Model (WSMX_O), v 0.1 available at <http://www.wsmo.org/2004/d13/d13.1/v0.1>

[WSMX Execution Semantics, 2004] E. Oren: WSMX Execution Semantics, WSMO Working Draft v01, 2004, available at <http://www.wsmo.org/2004/d13/d13.2/v0.1/20040531/index.pdf>

[WSMX Architecture, 2004] M. Moran, M. Zaremba: WSMX Architecture, WSMO Working Draft v01, 2004, available at <http://www.wsmo.org/2004/d13/d13.4/v0.1/20040622/>

[WSMX Implementation, 2004] M. Moran, E. Cimpian, A. Mocan, E. Oren: WSMX Implementation, WSMO Working Draft v01, 2004, available at <http://www.wsmo.org/2004/d13/d13.5/v0.1/20040719/>

Acknowledgments

The work is funded by the European Commission under the projects DIP, Knowledge Web, Ontoweb, SEKT, and SWWS; by Science Foundation Ireland under the DERI-Lion project; and by the Austrian government under the CoOperate programm.

The editors would like to thank to all the [members of the WSMO working group](#) for their advices and inputs to this document.



[webmaster](#)