

Reasoning and Change Management in Modular Ontologies

Heiner Stuckenschmidt ^{a,*} Michel Klein ^b

^a*Universität Mannheim*

^b*Vrije Universiteit Amsterdam, The Netherlands*

Abstract

The benefits of modular representations are well known from many areas of computer science. In this paper, we concentrate on the benefits of modular ontologies with respect to local containment of terminological reasoning. We define an architecture for modular ontologies that supports local reasoning by compiling implied subsumption relations. We further address the problem of guaranteeing the integrity of a modular ontology in the presence of local changes. We propose a strategy for analyzing changes and guiding the process of updating compiled information.

Key words: Ontologies, Reasoning, Distributed Knowledge Representation, Change Management

1 Motivation

Currently, research in the area of the semantic web is in a state where ontologies are ready to be applied in real applications such as semantic web portals, information retrieval or information integration. In order to lower the effort of building ontology-based applications, there is a clear need for a representational and computational infrastructure in terms of general purpose tools for building, storing and accessing ontologies. A number of such tools have been developed, i.e. ontology editors [1,2], reasoning systems [3,4] and more recently storage and query systems (e.g. [5]). Most of these tools, however,

* Corresponding author. Address: Universität Mannheim, Institut für Praktische Informatik, A5, 6, 68159 Mannheim, DE, email: heiner@informatik.uni-mannheim.de, Tel.: +49 621 181 2530

treat ontologies as monolithic entities and provide little support for specifying, storing and accessing ontologies in a modular manner. Existing proposals trying to fill this gap lack a formal underpinning.

1.1 Why Modularization ?

There are many reasons for thinking about ontology modularization. Our work is mainly driven by three arguments. These also bias the solution we propose, as it focusses on the following aspects.

Distributed Systems: In distributed environments like the semantic web, the question for modularization arises naturally. Ontologies in different places are built independent of each other and can be assumed to be highly heterogeneous. Unrestricted referencing to concepts in a remote ontology can therefore lead to serious semantic problems as the domain of interpretation may differ even if concepts appear to be the same on a conceptual level. The introduction of modules with local semantics can help to overcome this problem.

Large Ontologies: Modularization is not only desirable in distributed environments, it also helps to manage very large ontologies that we find in medicine or biology, for example. These ontologies, which sometimes contain more than a hundred thousand concepts, are hard to maintain as changes are not contained locally but can affect large parts of the model. Another argument for modularization in the presence of large ontologies is re-use: in most cases, we are not interested in the complete ontology when building a new system, but only in a specific part. Experiences from software engineering shows that modules provide a good level of abstraction to support maintenance and re-use.

Efficient Reasoning: A specific problem with distributed ontologies as well as with very large models is the efficiency of reasoning. While the pure size of the ontologies causes problems in the latter case, hidden dependencies and cyclic references can cause serious problems in a distributed setting. The introduction of modules with local semantics and clear interfaces will help to analyze distributed systems and provides a basis for the development of methods for localizing inference.

1.2 Requirements

There are a couple of requirements a modular ontology architecture has to fulfill in order to improve ontology maintenance and reasoning in the way suggested above. The requirements will be the main guidelines for the design of our solution proposed in this work.

Loose Coupling: In general, we cannot assume that two ontology modules have anything in common. This holds for the conceptualization as well as for the specific logical language used as well as for the interpretation of objects, concepts or relations. Our architecture has to reflect this by providing an extremely loose coupling of modules. In particular, we have to prevent unwanted interactions between modules. For this purpose, mappings between modules have to be distinguished from local definitions on the semantic as well as the conceptual level.

Self-Containment: In order to facilitate the re-use of individual modules from a larger, possibly interconnected system, we have to make sure that modules are self-contained. In particular, it should be possible to perform certain reasoning tasks such as subsumption or query answering within a single module without having to access other modules. This is also important if we want to provide efficient reasoning. Further we have to ensure correctness and whenever possible completeness of local reasoning for obvious reasons.

Integrity: The advantages of having self-contained ontology modules have their price in terms of potential inconsistencies that arise from changes in other ontology modules. While there is in our architecture no need to access other modules at reasoning time, the correctness of reasoning within a self contained module may still depend on knowledge in other ontologies. If this knowledge changes, reasoning results in a self-contained module may become incorrect with respect to the overall system, and we will not even notice it. We have to provide mechanisms for checking whether relevant knowledge in other systems has changed and for adapting the reasoning process if needed to ensure correctness.

1.3 Related Work

Our work relates to two main areas of research on representing and reasoning about ontological knowledge. The first is concerned with distributed and modular knowledge representation where we use ideas from theorem proving and knowledge engineering. The second area of related work is concerned with managing knowledge models. Here previous work exists in knowledge engineering as well as database information systems.

While the principle of modularity has widely been adopted in software engineering it has got less attention in the area of knowledge representation and reasoning. Some fundamental work on the modularization of representations can be found in the area of theorem proving. Farmer and colleagues promote the use of combinations of ‘Little Theories’, representations of a specific mathematical structure in order to reason about complex problems [6]. They

show the advantages of this modular approach in terms of reusability and reduced modeling effort. The idea of reusing and combining chunks of knowledge rather than building knowledge bases from scratch has later been adopted by the knowledge engineering community for building real-world knowledge bases (e.g. see [7]). McIlraith and Amir argue that a modularization of knowledge bases has also advantages for reasoning, even if the modularization is done a posteriori. They present algorithms for breaking down existing representations into a set of modules with minimal interaction and define reasoning procedures for propositional [8] and first-order logic [9]. The work reported is motivated by well established techniques from uncertain reasoning, where an a posteriori modularization of large theories is a common way to reduce runtime complexity (e.g. see [10]).

As we are interested in representations of ontological knowledge, approaches from the area of logics for representing terminologies, so-called description logics are of special interest for our work. In this area, we find the same arguments for a modularized representation as in the area of theorem proving. Rector proposes a strategy for modular implementation of ontologies using description logics [11]. The approach is based on a set of orthogonal taxonomies that provide a basis for defining more complex concepts. Rector argues for the benefits of this strategy in terms of easier creation and re-use of ontological knowledge. Buchheit and others propose a similar structuring on the language level by dividing the terminological part of a knowledge base into a schema part that corresponds to the basic taxonomies and a view part [12]. They show that this distinction can be used to achieve better run-time behavior for complex view languages. While these approaches still assume the overall model to be a single ontology providing a coherent conceptualization of the world, Giunchiglia and others propose a more radical approach to distributed representations. They propose the local model semantics as an extension of the standard semantics of first order logics [13]. This semantics allows different modules to represent different views on the same part of the world and the definition of directed partial mappings between different modules. Recently, Borgida and Serafini defined a distributed version of description logics based on local model semantics that has all advantages of the contextual representations [14].

As already mentioned, the problem of combining and reasoning with ontological modules has become of central importance in research on knowledge representation and reasoning on the semantic web. Standard languages for encoding ontological knowledge on the world wide web, i.e. the RDF schema [15] and the web ontology language OWL [16] provide some basic mechanisms for combining modular representations. The abilities to combine different models are restricted to the import of complete models and to the use of elements from

a different model in definitions by direct reference. It is assumed that references to external statements are only made for statements from imported models, however, this is strictly speaking not required. As a consequence, mappings rather implicitly exist in terms of mutual use of statements across models. Volz and colleagues discuss different interpretations of the import statement that range from purely syntactic to schema-aware interpretations of the imported knowledge [17]. An alternative way of relating different RDF models to each others that is much closer to our ideas is discussed by Oberle [18] who defines a view language for RDF and defines some consistency constraints for the resulting model.

1.4 *Our Approach*

In the following, we describe our approach to ontology modularization on an abstract level. We emphasize the main design decisions and motivate them on the basis of the requirements defined above. The technical details of the approach will be given in the subsequent sections.

View-Based Mappings: The first design decision concerns the way different ontology modules are connected. In our work, we adopt the approach of view-based information integration. In particular, ontology modules are connected by conjunctive queries and the extension of a concept in one module can be claimed to be equivalent to the (intentional) answer set of a conjunctive query over the vocabulary of another module. This way of connecting modules is more expressive than simple one-to-one mappings between concept names. Further, the same technique can be used to define relations of any arity based on other modules. Compared to the use of arbitrary axioms, our approach is less expressive. We decide to sacrifice a higher expressiveness for the sake of conceptual simplicity and desirable semantic properties such as directedness of the mapping.

Interface Compilation: The use of conjunctive queries guarantees a loose coupling on a conceptual and semantic level. However, it does not provide self-containment, because reasoning in an ontology module depends on the answer sets of the queries that are used to connect it to other modules. These answer sets have to be determined by actually querying the other ontology module. In order to make local reasoning independent from other modules, we use a knowledge compilation approach. The idea is to compute the result of each mapping query off-line and add the result as an axiom to the ontology module. At reasoning time these axioms replace the query, thus enabling local reasoning. As the results of queries are considered to be defined intentionally rather than extensionally, the result of the compilation of a query is not a set of instances retrieved from other modules, but a concept expression that contains all the information necessary to perform

local reasoning. In our case this expression is the conjunction of all concepts of the other ontology module that subsume the query expression.

Change Detection and Automatic Update: Our approach of compiling mappings and adding the result to the ontology models is very sensitive against changes in ontology modules. Once a query has been compiled, the correctness of reasoning can only be guaranteed as long as the concept hierarchy of the queried ontology module does not change. On the other hand, not every change in the hierarchy does really influence the compiled result. Problems only arise if concepts used in the query change or if the set of concepts subsuming the query is changed. In the second case, we will have to compile the interface again. In the first case we might even have to consider a redefinition of the query. In order to decide whether the compiled axiom is still valid, we propose a change detection mechanism that is based on a taxonomy of ontological changes and their impact on the concept hierarchy in combination with the position of the affected concept in the hierarchy. We further exploit an explicit representation of the dependencies between ontology modules in order to propagate changes in the system when necessary.

In the following, we first introduce a representational framework for modular ontologies that builds on top of existing work on distributed description logics (DDL) as a framework for reasoning about distributed ontologies and compare the resulting language with standard mechanisms for linking ontologies provided by the Web Ontology Language OWL. In Section 3 we define reasoning mechanisms for modular ontologies as a special case of general inference in distributed description logics. We further introduce the compilation of implied subsumption relations as a mechanisms for localizing reasoning and compare it with the distributed reasoning methods proposed for DDL. Section 4 discusses the problem of handling changes in external ontologies and their impact on compiled knowledge and proposes a heuristic for checking whether compiled knowledge has to be recomputed. We conclude with an example from a case study on ontology evolution in Section 5 and a discussion of the tradeoffs of our approach and possible extensions in Section 7.

2 Modular Ontologies

In this paper, we consider ontologies represented in the description logic \mathcal{SHIQ} . This choice is motivated by the fact that \mathcal{SHIQ} covers a large part of the expressive power of the Web Ontology Language OWL [16], more specifically of the language OWL-DL, a decidable sublanguage of OWL that directly corresponds to the logic \mathcal{SHOIQ} that extends \mathcal{SHIQ} with nominals [19]. We omit this extension in order to be able to base our framework on recent results on Distributed Description Logics [20,21] that provide us with basic mecha-

nisms for specifying links between concepts in different ontologies in a loose way. Before defining our notion of modular ontologies, we briefly introduce the logic \mathcal{SHIQ} as well as the basic notions of Distributed Description Logics. For background information about notation and naming in Description Logics, we refer to [22].

2.1 The \mathcal{SHIQ} Language

The basic modeling elements in Description Logics are concepts (classes of objects), roles (binary relations between objects) and individuals (named objects). Based on these modeling elements, Description Logics contain operators for specifying so-called concept expressions that can be used to specify necessary and sufficient conditions for membership in the concept they describe. Basic reasoning tasks associated with these kinds of logics are checking whether an expression is satisfiable (whether it is possible that an object satisfies the membership condition) and deciding subsumption between two concepts (deciding whether a concept expression implies another one). We now look at these issues on a more formal level.

Let \mathcal{C} be a set of concept names and \mathcal{R} a set of role names. Further let there be a set $R^+ \subseteq \mathcal{R}$ of transitive roles (i.e. for each $r \in R^+$ we have $r(x, y) \wedge r(y, z) \Rightarrow r(x, z)$). If now R^- denotes the inverse of a role (i.e. $r(x, y) \Rightarrow r^-(y, x)$) then we define the set of roles as $R \cup \{r^- | r \in R\}$. A role inclusion axiom is an expression $R \sqsubseteq S$ where R and S are roles. A role is called a simple role if it is not transitive and does not have transitive subroles with respect to the transitive closure of the role inclusion relation. The set of concepts (or concept expressions) in \mathcal{SHIQ} is the smallest set such that:

- \top and \perp are concept expressions
- every concept name A is a concept expression
- if C and D are concept expressions, r is a role, s is a simple role and n is a non-negative integer, then $\neg C$, $C \sqcap D$, $C \sqcup D$, $\forall r.C$, $\exists r.C$, $\geq nr.C$ and $\leq nr.C$ are concept expressions.

A general concept inclusion axiom is an expression $C \sqsubseteq D$ where C and D are concepts. A terminology is a set of general concept inclusion and role inclusion axioms.

The semantics of \mathcal{SHIQ} is defined in terms of an interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ where $\cdot^{\mathcal{I}}$ is a function that maps every concept on a subset of $\Delta^{\mathcal{I}}$ and every role on a subset of $\Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$ such that for all concepts C and D and for roles r where $\#M$ denotes the cardinality of M and $(r^{\mathcal{I}})^+$ the transitive closure of $r^{\mathcal{I}}$ we have:

- $\top^{\mathcal{I}} = \Delta^{\mathcal{I}}$ and $\perp^{\mathcal{I}} = \emptyset$
- $r^{\mathcal{I}} = (r^{\mathcal{I}})^+$ for $r \in R^+$ and $r^- = \{(y, x) \mid (x, y) \in r^{\mathcal{I}}\}$
- $(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} - C^{\mathcal{I}}$, $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$ and $(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}}$
- $(\forall r.C)^{\mathcal{I}} = \{x \mid \forall y.(x, y) \in r^{\mathcal{I}} \Rightarrow y \in C^{\mathcal{I}}\}$
- $(\exists r.C)^{\mathcal{I}} = \{x \mid \exists y.(x, y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\}$
- $(\geq n r.C)^{\mathcal{I}} = \{x \mid \#\{y.(x, y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\} \geq n\}$
- $(\leq n r.C)^{\mathcal{I}} = \{x \mid \#\{y.(x, y) \in r^{\mathcal{I}} \wedge y \in C^{\mathcal{I}}\} \leq n\}$

An interpretation satisfies a terminology \mathcal{T} if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for all general concept inclusions $C \sqsubseteq D$ in \mathcal{T} and $r^{\mathcal{I}} \subseteq s^{\mathcal{I}}$ for all role inclusion axioms $r \sqsubseteq s$ in \mathcal{T} . In this case we call \mathcal{I} a model for \mathcal{T} . A concept D subsumes a concept C in \mathcal{T} if $C \sqsubseteq D$ holds for all models of \mathcal{T} . In the remainder of the paper we will focus on the task of deciding whether a concept subsumes another one.

2.2 Distributed Description Logic

Distributed Description Logics as proposed in [20] provide a language for talking over sets of terminologies. For this purpose DDLs provide mechanisms for referring to terminologies and for defining rules that connect concepts in different terminologies. On the semantic level, DDLs extend the notion of interpretation introduced above to fit the distributed nature of the model and to reason about concept subsumption across terminologies.

Let I be a non-empty set of indices and $\{\mathcal{T}_i\}_{i \in I}$ a set of terminologies. We prefix inclusion axioms with the index of the terminology they belong to (i.e. $i : C$ denotes a concept in terminology \mathcal{T}_i and $j : C \sqsubseteq D$ a concept inclusion axioms from terminology \mathcal{T}_j). Note that $i : C$ and $j : C$ are different concepts. Semantic relations between concepts in different terminologies are represented in terms of axioms of the following form, where C and D are concepts in terminologies \mathcal{T}_i and \mathcal{T}_j , respectively:

- $i : C \xrightarrow{\sqsubseteq} j : D$ (into-rule)
- $i : C \xrightarrow{\supseteq} j : D$ (onto-rule)

These axioms are also called *bridge-rules*. An additional rule $i : C \xrightarrow{\equiv} j : D$ is defined as the conjunction of the two rules above. A distributed terminology \mathfrak{T} is now defined as a pair $(\{\mathcal{T}_i\}_{i \in I}, \{\mathfrak{B}_{ij}\}_{i \neq j \in I})$ where $\{\mathcal{T}_i\}_{i \in I}$ is a set of terminologies and $\{\mathfrak{B}_{ij}\}_{i \neq j \in I}$ is a set of bridge rules between these terminologies.

The semantics of distributed description logics is defined in terms of a global interpretation $\mathfrak{I} = (\{\mathcal{I}_i\}_{i \in I}, \{r_{ij}\}_{i \neq j \in I})$ where \mathcal{I}_i is an interpretation for terminology \mathcal{T}_i as defined above and $r_{ij} \subseteq \Delta^{\mathcal{I}_i} \times \Delta^{\mathcal{I}_j}$ is a domain relation connecting elements of the interpretation domains of terminologies \mathcal{T}_i and \mathcal{T}_j . We use

$r_{ij}(x)$ to denote $\{y \in \Delta^{\mathcal{I}_j} \mid (x, y) \in r_{ij}\}$ and $r_{ij}(C)$ to denote $\bigcup_{x \in C} r_{ij}(x)$.

A distributed interpretation \mathfrak{I} satisfies a distributed terminology \mathfrak{T} if:

- \mathcal{I}_i satisfies \mathcal{T}_i for all $i \in I$
- $r_{ij}(C^{\mathcal{I}_i}) \subseteq D^{\mathcal{I}_j}$ for all $i : C \xrightarrow{\sqsubseteq} j : D$ in \mathfrak{B}_{ij}
- $r_{ij}(C^{\mathcal{I}_i}) \supseteq D^{\mathcal{I}_j}$ for all $i : C \xrightarrow{\sqsupseteq} j : D$ in \mathfrak{B}_{ij}

In this case we call \mathfrak{I} a model for \mathfrak{T} . A concept $i : D$ subsumes a concept $i : C$ ($i : C \sqsubseteq D$) if for all models of \mathfrak{T} we have $C^{\mathcal{I}_i} \subseteq D^{\mathcal{I}_i}$.

2.3 Modular Ontologies

We can now define our notion of a modular ontology in terms of Distributed Description Logics. In fact, our notion of a modular ontology is a restricted form of distributed terminology as defined above. The restrictions we impose concern the architecture of the distributed terminology as well as the expressiveness of semantic relations between terminologies. These restrictions are motivated by the aims of (1) providing an alternative to the standard notion of import in OWL and (2) the goal of providing support for localized reasoning and maintenance of the modular ontology. In the following, we will first discuss the architecture of a modular ontology and then introduce the restrictions imposed on semantic relations.

2.3.1 Architecture

As described above, DDLs make a clear distinction between terminologies and semantic mappings between them in terms of bridge rules, which in principle are independent of the terminologies. This makes the model quite flexible; for example, it permits to have different sets of mapping rules connecting the same set of ontologies. In this way it is possible to encode different views on how the terminologies relate to each other. In contrary to that, our aim is to resemble the use of external knowledge in a terminology similar to the ability of OWL to use concept and role names defined in different terminologies. This view is different from the model of Distributed Description Logics as it makes the semantic links to other models part of the terminology. Being part of the terminology implies that there is only one way of connecting to these external definitions which is assumed to be agreed on by the users of the local terminology.

We achieve this localization of semantic relations by inventing the notion of externally defined concepts in a terminology. We divide the set of concept names in a terminology into internally defined concepts \mathcal{C}_I and externally

defined concepts \mathcal{C}_E resulting into the following definition of the set of all concept names \mathcal{C} :

$$\mathcal{C} = \mathcal{C}_I \cup \mathcal{C}_E, \mathcal{C}_I \cap \mathcal{C}_E = \emptyset \quad (1)$$

We consider externally defined concepts to be concept names linked to a concept expression defined in another terminology using bridge rules. An external concept definition in terminology \mathcal{T}_i is an axiom of the form: $i : C \equiv \mathcal{T}_j : D$ where C is a concept name in \mathcal{T}_i , \mathcal{T}_j is a terminology different from the one in which the external concept is defined and D is a concept expression in \mathcal{T}_j . Note that although D is syntactically represented in \mathcal{T}_i it actually represents a concept in \mathcal{T}_j . In particular the expression D is only allowed to contain concepts defined in \mathcal{T}_j . This definition is very close to the OWL mechanism of using concept and role names from other name spaces in definitions.

We give external concept definitions a semantics in terms of distributed description logics by defining external concept definitions to be an alternative notation for a pair of bridge rules:

$$i : C \equiv \mathcal{T}_j : D \Leftrightarrow j : D \xrightarrow{\sqsubseteq} i : C \wedge j : D \xrightarrow{\sqsupseteq} i : C$$

Semantically, this is equivalent to

$$C_i^{\mathcal{T}} = r_{ji}(D_j^{\mathcal{T}}) \quad (2)$$

The correspondence between external concept definitions and bridge rules allows us to base our further investigations on the formal results that have been established for distributed *SHIQ* terminologies.

2.3.2 Restricting Mapping Expressiveness

In Distributed Description Logics, there are no restrictions on the antecedent of a bridge rule—except that it has to be a valid concept of the source terminology. In our framework, we restrict this freedom for the sake of an easier maintenance of the semantic relations between terminologies. **HEINER:check meaning of sentence:** We require the external concept definition to be equivalent to a conjunctive query of arity one over concepts and roles in the external terminology. The rationale for this choice is three-fold. First of all, work on schema integration has shown that conjunctive queries are well suited for describing complex mappings between conceptual schema (see for example [23]). Secondly, as we will see in section 4 conjunctive queries show a useful behavior with respect to predicting the impact of changes in the external terminology on the local one. Besides this, a corresponding query

language for the Web Ontology Language OWL is currently being developed under the name OWL-QL [24]. Basing our mappings on conjunctive queries will enable us to use this emerging standard for representing mappings between terminologies on the syntactic level. Finally, there are methods for answering conjunctive queries over terminologies as well as for deciding subsumption between conjunctive queries over terminologies that turn out to be useful in the context of our work.

Horrocks and Tessaris describe an approach for translating conjunctive queries into an equivalent concept expression [25]. The concept expression can be classified into the knowledge base and standard DL inference methods can be used to check whether an object is an instance of the concept corresponding to the query expression or whether a query subsumes another one. This approach makes use of the fact that binary relations in a conjunctive query can be translated into existential restrictions in such a way that logical consequence is preserved. A consequence of this result is that we can translate conjunctive queries into concept expressions over the sublanguage of \mathcal{SHIQ} containing the following concept building operators (compare section 2.1):

$$C, D \longrightarrow \top \mid \perp \mid A \mid C \sqcap D \mid \exists R.C \mid \exists R^-.C$$

In order to restrict the semantic correspondences between terminologies in our model, we now only allow the concept expressions D in the definition of external concepts to be valid concepts over terminology \mathcal{T}_j with respect to the sublanguage defined above. We denote such concepts as D_Q and consider external concept expressions of the form $i : C \equiv j : D_Q$.

3 Reasoning in Modular Ontologies

The direct correspondence of our framework to Distributed Description Logics allows us to base inference in modular ontologies on known results for the corresponding logics. In particular, we can provide completeness and complexity for reasoning in modular ontologies. We extend the existing work on reasoning in DDLs with the notion of compilation of implied subsumption relations. Specifically, we use reasoning methods for Distributed Description Logics to derive subsumption relations between externally defined concepts in modules and explicitly add the derived subsumption relations as axioms to the module. The results of [21] guarantee that after this compilation step reasoning can be performed locally without considering other modules unless there are changes in the system.

In the following, we first briefly review basic definitions of reasoning in distributed description logics and prove that it has the same worst-case complexity

as reasoning in the local logics used. We then present the compilation of subsumption relations and discuss conditions for completeness and consistency.

3.1 Reasoning based on DDLs

Reasoning in DDLs differs from reasoning in traditional Description Logics by the way knowledge is propagated between T-Boxes by certain combinations of bridge rules. The simplest case in which knowledge is propagated is the following:

$$\frac{i : A \xrightarrow{\exists} j : G, i : B \xrightarrow{\sqsubseteq} j : H, i : A \sqsubseteq B}{j : G \sqsubseteq H} \quad (3)$$

This means that the subsumption between two concepts in a T-Box can depend on the subsumption between two concepts in a different T-Box if the subsumed concepts are linked by the *onto*- and the subsuming concepts by an *into* rule. In languages that support disjunction, this basic propagation rule can be generalized to subsumption between a concept and a disjunction of other concepts in the following way:

$$\frac{i : A \xrightarrow{\exists} j : G, i : B_k \xrightarrow{\sqsubseteq} j : H_k (1 \leq k \leq n), i : A \sqsubseteq \bigsqcup_{k=1}^n B}{j : G \sqsubseteq \bigsqcup_{k=1}^n H_k} \quad (4)$$

It has been shown that this general propagation rule completely describes reasoning in DDLs that goes beyond well known methods for reasoning in Description Logics. To be more specific, adding the inference rule in equation 4 to existing tableaux reasoning methods lead to a correct and complete method for reasoning in DDLs. A corresponding result using a fixed point operator is given in [21]. Based on these results, we can define a general inference rule for the case of modular ontologies in the following way:

$$\frac{i : A \equiv j : G, i : B_k \equiv j : H_k (1 \leq k \leq n), i : A \sqsubseteq \bigsqcup_{k=1}^n B}{j : G \sqsubseteq \bigsqcup_{k=1}^n H_k} \quad (5)$$

There are a number of consequences of this result for reasoning in modular ontologies.

Correctness and Completeness From the basic propagation rule, we can see that subsumption between externally defined concepts follows from subsumption of their definitions in the (same) external module. This is because each external concept definition corresponds to an *into* and an *onto* rule between the concept name and its definition. The language we consider is \mathcal{SHIQ} and therefore we have to consider the general propagation rule because we have disjunction in our language. This means that it is not enough to simply check whether subsumption between the definitions of two externally defined concepts in the external module is complete, but we have to consider all subsets of the set of external concepts. We will discuss this point in more detail in the next section.

Complexity As we reduce reasoning in modular ontologies to reasoning in DDLs with \mathcal{SHIQ} as a local language, complexity results can be derived from known results on reasoning in \mathcal{SHIQ} and Distributed Description Logics.

Theorem 1 (Complexity) *Reasoning in modular ontologies is Exp-Time Complete.*

Proof 1 *We show that reasoning in modular ontologies has the same complexity as reasoning in the local language. As reasoning in \mathcal{SHIQ} is Exp-Time Complete [26] this establishes the result.*

- *Reasoning in modular ontologies is at least as hard as reasoning in \mathcal{SHIQ} : in the extreme case a modular ontology consists only of a single module without external concepts, thus reasoning in modular ontologies is equivalent to reasoning in \mathcal{SHIQ}*
- *Reasoning in modular ontology is not harder than reasoning in \mathcal{SHIQ} : we reduce reasoning in modular ontologies to reasoning in DDLs. There exists a reduction of reasoning in DDLs with \mathcal{SHIQ} as a local language to \mathcal{SHIQ} . Both reductions are linear in the size of the resulting terminology and therefore do not change the complexity class.*

This result shows that the complexity of reasoning in modular ontologies is not worse than reasoning in the web ontology language. Using the reduction of DDLs to \mathcal{SHIQ} it is even possible to use existing OWL reasoners for reasoning with modular ontologies. Although practical implementations of OWL reasoners have shown that good average case performance can be achieved, the worst case complexity is still very high and asks for further optimization.

3.2 *Compilation and Integrity*

Existing reasoners for expressive Description Logics are highly optimized with respect to the deciding subsumption in the context of a single T-Box. Serafini and Tamilin present a distributed reasoning system that extends existing reasoners to distributed T-Boxes [27]. In theory, this system is complete with respect to the propagation rules described above and has—as we have argued—the same worst-case complexity. In practice, however, reasoning with multiple, possible distributed modules, brings some new problems with respect to completeness and reasoning performance. First of all the completeness of the distributed reasoners depends on the availability of local reasoners for all T-Boxes in the system. In a loosely coupled network without central control this cannot always be guaranteed as network nodes can be unreachable or even leave the network. In this case, necessary subsumption tests cannot be performed at these nodes leading to a possible incompleteness. Another problem currently not addressed in the work of Serafini and Tamilin are performance problems due to communication costs between the different nodes in the system. Work in the area of distributed databases has shown that communication costs often become serious bottlenecks in distributed systems.

In order to overcome these problems we propose to compute subsumption relations between external concepts offline and store them as explicit axioms in the local ontologies. If we compute these relations using the reasoner mentioned above we have the guarantee that reasoning about subsumption in each module can be done without caring about the availability of other nodes in the network. This also has the advantage that no communication costs occur as part of online reasoning.

Of course these runtime benefits have their price in terms of computational complexity of the compilation step. The completeness of the propagation rule in expression 4 tells us that to be independent from other modules we only have to consider subsumption relations between externally defined concepts, as only such subsumption relations can be propagated from outside. What we have to check is subsumption between each external concept and the disjunction of all combinations of other external concepts. For a local module, this process is defined in algorithm 1.

If we denote the number of external concepts \mathcal{C}_E as n , the worst-time complexity of the compilation method is $O(n \cdot 2^{(n-1)})$ as can easily be seen from the algorithm. As deciding the subsumption relation in the conditional statement of the algorithm itself is already Exp-Time Complete and this test has to be carried out an exponential number of times with respect to the number of external concepts, compiling all implied statements is computationally very expensive. We therefore do not want to perform the compilation step

Algorithm 1 compile

Require: An T-Box \mathcal{T} with external concepts \mathcal{C}_E

```
for all  $c \in \mathcal{C}_E$  do  
  candidates :=  $\mathcal{P}(\mathcal{C}_E - \{c\})$   
  for all  $d \in$  candidates do  
    if  $\mathcal{T} \models c \sqsubseteq \bigsqcup_{e \in d} e$  then  
       $\mathcal{T} := \mathcal{T} \cup \{c \sqsubseteq \bigsqcup_{e \in d} e\}$   
    end if  
  end for  
end for
```

more often than absolutely necessary to guarantee that local reasoning is still complete.

While the results of Serafini and others guarantee that local reasoning is correct and complete at the time the compilation is carried out, a problem occurs when changes are made to the system. Changes in the definitions of the external concepts, but also changes in the definitions of concepts and relations in other modules can make local reasoning incomplete or inconsistent. In order to prevent situations in which local reasoning is not correct and complete any more we introduce the notion of integrity of a modular ontology.

Definition 1 (Integrity) *Let $\mathfrak{T} = (\{\mathcal{T}_i\}_{i \in I}, \{\mathfrak{B}_{ij}\}_{i \neq j \in I})$ be a modular ontology with interpretation $\mathfrak{J} = (\{\mathcal{I}_i\}_{i \in I}, \{r_{ij}\}_{i \neq j \in I})$ then we say that integrity holds for \mathfrak{T} if for all $\mathcal{T}'_i = \text{compile}(\mathcal{T}_i)$ with interpretation \mathcal{I}'_i we have:*

$$\mathcal{I}'_i \models C \sqsubseteq D \Leftrightarrow \mathfrak{J} \models i : C \sqsubseteq D$$

for any pair of legal concept expressions C and D in \mathcal{T}'_i .

The notion of integrity gives us a criterion for deciding whether compiled results are still valid. What the definition does not provide is an operational account for checking it. A direct use of the definition would involve a complete check of all derivable subsumption relations. As we have argued above this approach is extremely expensive. In the following, we therefore present a heuristic approach for checking integrity in modular ontologies that is driven by changes made to the ontology. The approach is able of determining situations in which changes to a modular ontology do not affect integrity and therefore no re-compilation is necessary.

4 Evolution Management

As we have argued above, guaranteeing integrity of compiled subsumption relations is the main problem in modular ontologies. In principle, all compiled subsumption relations have to be recomputed to test whether they are still valid in the given state of the system. This of course means sacrificing the advantages of the compilation approach in terms of local reasoning and reduced complexity. Fortunately, we can do better than checking all compiled axioms each time we perform reasoning.

The first possible improvement is to move away from an active checking for changes towards a mechanism where each local module remembers and records changes made to it. We can also think of a system where individual modules actively notify other modules of changes to its local knowledge. This frees us from doing a complete check of the compiled knowledge and allows us to concentrate on these parts of the knowledge that actually were subject to changes.

The second improvement is in terms of an analysis of the impact a change in another module actually has on compiled knowledge. This is important as in existing scenarios it turns out that a large part of the changes do not really affect the logical theory but are rather changes to the syntactic representation or changes in the naming of concepts and relations. While the latter have to be propagated to the definitions of the external concepts, they do not actually affect the compiled subsumption relations. Further, even if the logical theory underlying the ontology is affected by a change, this does not mean that it affects the compiled subsumption relations. This means that we have to find ways to distinguish changes that do have an impact on compiled relations from those that do not have an impact. In fact, the choice to restrict the language admissible in the definitions of external concepts allows us to precisely characterize these kinds of changes.

In the following, we concentrate on the analysis of the impact of changes on the validity of compiled subsumption relations. We first give a characterization of harmless (changes that do not have an effect on compiled subsumption relations) and harmful (changes that do have a potential effect on compiled subsumption relations) changes. We then present mechanisms for classifying changes as harmless or harmful based on a syntactic analysis of changes made to an ontology. Finally, we present a simple mechanism that uses this information to decide whether the knowledge in a module has to be re-compiled.

4.1 Determining harmless changes

As compiled knowledge reflects subsumption relations between a query concept and a disjunction of other query concepts a harmless change is a set of modifications to an ontology that does not change these subsumption relations. Finding harmless changes is therefore a matter of deciding whether the modifications affect the subsumption relation between a query concept and a disjunction of other query concepts. It is quite obvious that a complete decision procedure for this problem has the same complexity as general subsumption reasoning in the modular ontology and does therefore not improve the situation. For this reason, we propose a sound but incomplete method that abstracts from the detailed definition of concepts and uses the semantic relation between the old and the new version of a concept in the following way.

The method considers every concept and relation in an ontology module that has been subject to a change. Assuming that C represents the concept under consideration before and C' the concept after the change, there are four ways in which the old version C may relate to the new version C' :

- (1) the meaning of a concept is not changed: $C \equiv C'$ (e.g. because the change was in another part of the ontology, or because it was only syntactical);
- (2) the meaning of a concept is changed in such a way that concept becomes more general: $C \sqsubseteq C'$;
- (3) the meaning of a concept is changed in such a way that concept becomes more specific: $C' \sqsubseteq C$;
- (4) the meaning of a concept is changed in such a way that there is no subsumption relationship between C and C' .

The same is true for relations that are subject to change. The next question is how these different types of changes influence the interpretation of query concepts. We take advantage of the fact that there is a very tight relation between changes in concepts of the external ontology and implied changes to the query concepts using these concepts:

Lemma 1 (monotonicity of effect) *Let $C \equiv T_j : D$ an external concept expression. let further be $c(D)$ be the set of all concept names and $r(D)$ the set of all relation names occurring in D .*

- C' is more general than C if there is an $x \in c(D) \cup r(D)$ such that x' is more general than x .
- C' is more specific than C if there is an $x \in c(D) \cup r(D)$ such that x' is more specific than x .

Note that the implication does not hold in the other direction.

We can exploit this relation between the interpretation of external concepts and the concept names in their definitions in order to identify the effect of changes in the external ontology on the subsumption relations between different query concepts. First of all, the above result directly generalizes to multiple changes with the same effect, i.e. a query Q becomes more general (specific) or stays the same if none of the elements in $c(Q) \cup r(Q)$ become more specific (general). Further, the subsumption relation between an external concept C and the disjunction of other external concepts does not change if all concepts in the disjunction become more general or if the concept C becomes more specific. Combining these two observations, we derive the following characterization of harmless change.

Theorem 2 (harmless change) *Let $C_0 \equiv \mathcal{T}_j : D_0, C_1 \equiv \mathcal{T}_j : D_1, \dots, C_m \equiv \mathcal{T}_j : D_m$ be external concept definitions such that $\mathfrak{I} \models C_0 \sqsubseteq C_1 \sqcup \dots \sqcup C_m$, then a change is harmless with respect to the subsumption relation above if:*

- $X' \sqsubseteq X$ for all $X \in c(D_0) \cup r(D_0)$,
- $X' \sqsupseteq X$ for all $X \in c(D_i) \cup r(D_i)$, $i = 1, \dots, m$

Note again that the implication does not hold in the opposite direction.

The theorem provides us with a correct but incomplete method for deciding whether a change is harmless given that we know the semantic relation between the old and the new definition of concepts and relations that were subject to changes. This method is a very basic version of the underlying idea of assessing the impact of changes. We can think of more complete versions of the method that uses a deeper analysis of the structure of the concept expressions involved. Our experiences are, however, that this basic heuristic already covers most cases that occur in practice, especially, because the definition above includes cases where most of the concepts stay unchanged.

4.2 Characterizing changes

Now that we are able to determine the consequence of changes in the concept hierarchy on the integrity of the mapping, we still need to know what the effect of specific modifications on the interpretation of a concept is (i.e. whether it becomes more general or more specific). As our goal is to determine the integrity of mappings without having to do classification, we describe what theoretically could happen to a concept as result of a modification in the ontology. To do so, we have listed all possible change operations to an ontology according to the OWL ¹ knowledge model in the same style as done in [28].

¹ See <http://www.w3.org/TR/owl-features/>

The list of operations is in principle extendable to other knowledge models.

The list of change operations consists of two types of operations: (1) *atomic change operations*, such as *add range restriction* or *delete subconcept relation* and (2) *complex change operations*, which consist of multiple atomic operations and/or incorporate some additional knowledge. Complex changes are often more useful to specify effects than the atomic changes, as incorporate some of the semantic consequences. For example, for operations like *concept moved up*, or *domain enlarged*, we can specify the effect more accurately than for the atomic operations *superconcept changed* and *domain modified*.² Atomic changes can be detected at a structural level, i.e. by comparing the old and new definition of a concept, and are therefore computationally cheap with a linear complexity. To identify complex changes, we also need to take some of the semantic relations in the ontology into account. This makes the complexity of the identification of complex changes potentially as bad as determining subsumption in *SHIQ*, i.e. Exp-Time Complete. However, in practice many complex changes can be detected at a structural level, e.g. by looking at explicitly stated subclass relations.

Table 1 contains some examples of operations and their effect on the classification of concepts. The table only shows a few examples, although our full ontology of change operations contains around 120 operations. This number is not fixed, as new complex changes can be defined. A snapshot of the change ontology can be found online³.

The specification of effects is not complete, in the sense that it describes “worst-case” scenarios, and that for some operations the effect is “unknown” (i.e. unpredictable). In contrast to [30] who provide complete semantics of changes, we prefer to use heuristics in order to avoid expensive reasoning about the impact of changes. By restricting the change detection to changes that can be detected at a structural level, the complexity of our change detection heuristic is linear.

4.3 Update management

With the elements that we described in this section, we now have a complete procedure to determine whether compiled knowledge in an ontology module is still valid when the external ontology modules are changed. The complete

² For a complete list, see [29].

³ <http://ontoview.org/changes/2/1>

Table 1

Some modifications to an ontology and their effects on the classification of concepts in the hierarchy.

Operation	Effect
Add a role restriction to concept C	C : Specialized
<i>Complex</i> : change the superconcept of concept C to a concept lower in the hierarchy	C : Specialized
<i>Complex</i> : restrict the range of a role R (<i>effect on all C that have a restriction on R</i>)	R : Specialized, C : Specialized
Remove a superconcept relation of a concept C	C : Generalized
Change the concept definition of C from primitive to defined	C : Generalized
Add a concept definition A	C : Unknown
<i>Complex</i> : add a (not further specified) subconcept A of C	C : No effect
Define a role R as functional	R : Specialized

procedure is as follows. For each external concept C :

- (1) determine the changes that are performed in the external ontology (e.g. by using the record of changes);
- (2) heuristically determine the effect of the changes on the interpretation of the concepts and relations (where multiple changes to a concept or relation that have an opposite effect lead to the effect “unknown”);
- (3) create a list of all concept and relation names that occur in the external concept expression D ;
- (4) check whether all concepts and relation in this list are unchanged or became more specific;
- (5) create a list of all concept and relation names that occur in external concepts expressions other than D ;
- (6) check whether all concepts and relation in this list are unchanged or became more general.

In cases where we cannot guarantee that integrity is preserved, we recompute and re-compile the implied subsumption statements. We thus restore integrity and make correct local reasoning possible.

We describe the procedure in a more structured way in Algorithm 2. The algorithm triggers a (re-)compilation step only if it is required in order to re-sume integrity. Otherwise no action is taken, because the previously compiled knowledge is still valid. In principle, all the steps can be automated. A tool that helps to automate steps 1 and 2 is described in [31]. This tool will compare two versions of an ontology and derive the list of change operations that is necessary to transform the one into the other.

Algorithm 2 Update

Require: Ontology Module M

Require: Ontology Module M_j

```
for all compiled axioms  $C_1 \sqsubseteq D_1 \sqcup \dots \sqcup D_m$  in  $M^c$  do
  for all  $X \in c(C) \cup r(C)$  do
    if effect on  $X$  is 'generalized' or 'unknown' then
       $M^c := Recompile(M, M_j)$ 
    end if
  end for
  for all  $D_i, i = 1, \dots, m$  do
    for all  $X \in c(D_i) \cup r(D_i)$  do
      if effect on  $X$  is 'specialized' or 'unknown' then
         $M^c := Recompile(M, M_j)$ 
      end if
    end for
  end for
end for
```

The worst-time complexity of the this update procedure once the effects of changes are known is linear in the number of concept and relation names occurring in compiled subsumption statements (in the worst case the effect of a change on every concept and relation name in the compiled axioms has to be checked) . In practice, we expect the number of compiled axioms to be relatively small leading to a quite efficient procedure.

5 Application in a case study

In order to support the claims made about the advantage of modular ontologies, we applied our model in a small case study that has been carried out in the course of the WonderWeb project⁴. Our main intention was to show that the update-management procedure presented in the last section can be used to avoid the computation of subsumption relations in many cases. For this purpose, we defined a small example ontology using mappings to an ontology in the human resource (HR) domain. We used the changes that occurred in the HR ontology during the different steps of the case study and determined the impact on our example ontology. Besides this, the case study provides us with examples of implied subsumption some of which are non-trivial but likely to occur in real-life situations.

⁴ See <http://wonderweb.semanticweb.org>.

5.1 The WonderWeb case study

WonderWeb is a EU/IST project that ran from 2002 till 2004. The aim of the project was to develop and demonstrate the infrastructure required for the large-scale deployment of ontologies as the foundation for the Semantic Web. In the context of this project the *Descriptive Ontology for Linguistic and Cognitive Engineering* (DOLCE) [32] has been developed. DOLCE is the first module in a library of foundational ontologies. The taxonomy of the most basic categories assumed in DOLCE consists of concepts such as “location”, “social agent”, event” and “region”. One of the roles of foundational ontologies in general and DOLCE in particular is to clarify hidden assumptions underlying existing ontologies or linguistic resources by manually mapping existing categories into the categories defined in the foundational ontology [33]. Inconsistencies in the combined ontology point to modeling errors or wrong assumptions in the original ontology. Solving the inconsistencies will improve the quality of the initial ontology. The methodology around DOLCE describes a three step procedure for mapping existing ontologies to DOLCE: alignment, refinement and tidying.

The case study that has been carried out in the project integrates different methods in the life-cycle of an ontology that is used on the Semantic Web, i.e. ontology creation, ontology refinement and ontology deployment. The case study starts from an existing database schema in the human resource (HR) domain. A first version of an ontology is created by a tool that automatically converts a schema into an ontology [34]. In the next phase, the quality of the ontology is improved by manually mapping this ontology to DOLCE. First, the HR ontology is aligned with DOLCE, and in several successive steps the resulting ontology is further refined. During this process, the ontology changes continuously, which causes problems when other ontologies refer to definitions in the evolving ontology. Therefore, in our case study, evolution management is important during the entire life-cycle of the ontology-development process.

The original HR ontology combined with DOLCE is referred to as the DOLCE+HR ontology. In order to demonstrate the update management procedure, we created another ontology (which we call the *local ontology*) that uses terms and definitions from the evolving DOLCE+HR ontology (the *external ontology*). The local ontology defines the concept *FulltimeEmployee* with a superconcept *Employee* and two subconcepts *DepartmentMember* and *HeadOfDepartment*, using terms from the DOLCE+HR ontology.

The specific problem in our case is that the changes in the DOLCE+HR ontology could affect the reasoning in the local ontology. We want to be able

to predict whether or not the reasoning in the local ontology is still valid for specific changes in the external ontology.

The evolution of the DOLCE+HR ontology consisted of several steps. Each of these steps involves some typical changes. We will briefly summarize them and show some changes that are typical for a specific step.

- In the first step, the extracted HR ontology is aligned with the DOLCE foundational ontology, i.e. the concepts and roles in the HR ontology are connected to concepts and roles in the DOLCE ontology via subsumption relations. For example, the concept *Departments* from the HR ontology is made a subconcept of *Social-Unit* in DOLCE.
- The refinement step involves a lot of changes. Some role restrictions are added, and some additional concepts and roles are created to define the HR concepts more precisely. For example, the concept *Administrative-Unit* is introduced as a new subconcept of *Social-Unit*, and the concept *Departments* is made a subconcept of it. Also, the range of the role *email* is restricted from *Abstract-Region* to its new subconcept *Email*.
- In the next step, a number of concepts and roles are renamed to names that better reflect their meaning. For example, *Departments* is renamed to *Department* (singular), and the two different variants of the role *manager-id* are renamed to *employee-manager* and *department-manager*.
- In the final step, the tidying step, all roles and concepts that are not necessary any more are removed and transformed into role restrictions. For example, the role *employee-email* is deleted and replaced by an existential restriction in the concept *Employee* on the role *abstract-location* to the concept *Email*.

5.2 Modularization in the case study

If we now consider the modularization in the case study, we have a local ontology with a concept hierarchy that is built up by the following explicitly stated subsumption relations (see Fig. ?? again):

$$\begin{aligned}
 & FulltimeEmployee \sqsubseteq Employee \\
 & DepartmentMember \sqsubseteq FulltimeEmployee \\
 & HeadOfDepartment \sqsubseteq FulltimeEmployee
 \end{aligned}$$

This ontology introduces *FulltimeEmployee* as a new concept, not present in the case study ontology. Consequently, this concept is only defined in terms

of its relation to other concepts in the local ontology.

All other concepts are externally defined in terms of ontology based queries over the case study ontology. The first external definition concerns the concept *Employee* that is equivalent to the “Employee” concept in the case study ontology. This can be defined by the following trivial view:

$$Employee \equiv HR : Employee(x)$$

Another concept that is externally defined is the “head of department” concept. We define it to be the set of all instances that are in the range of the “department manager” role. The definition of this view given below shows that our approach is flexible enough to define concepts in terms of relations.

$$HeadOfDepartment \equiv HR : \exists y [departmentManager(y, x)]$$

An example of a more complex external concept definition is the concept *DepartmentMember*, which is defined using a query that consists of three conjuncts, claiming that a department member is an employee that is in the “has-member” role with a department.

$$DepartmentMember \equiv HR : \exists y [Department(y) \wedge hasmember(y, x) \wedge Employee(x)] \quad (6)$$

5.2.1 Implied subsumption relations

If we now consider logical reasoning about these external definitions, we immediately see that the definition of employee subsumes the definition of *DepartmentMember*, as the former occurs as part of the definition of the latter.

$$\models DepartmentMember \sqsubseteq Employee \quad (7)$$

At a first glance, there is no relation between the definition of a head of department and the two other statements as it does not use any of the concept or role names. However, when we use the background knowledge provided by the case study ontology we can derive some implied subsumption relations. The reasoning is as follows. Because the range of “department manager” is

set to “department” and the domain to “employee”, the definition of “Head-ofDepartment” is equivalent to:

$$\exists y[Department(y) \wedge departmentManager(y, x) \wedge Employee(x)]$$

As we further know that manager is a subconcept of employee and “departmentManager” is a subrole of “has-member”, we can derive the following subsumption relation between the externally defined concepts:

$$\models HeadOfDepartment \sqsubseteq Employee \quad (8)$$

$$\models HeadOfDepartment \sqsubseteq DepartmentMember \quad (9)$$

When the relations 7–9 are added to the local ontology, it possible to do subsumption reasoning without having to access the DOLCE+HR ontology any more.

5.3 Updating the models

We will now illustrate that the conclusions of the procedure are correct by studying the impact of changes mentioned in the problem statement.

5.3.1 Example 1: the employee concept

The first change we observed is the removal of relations from the employee concept. Our rules tell us that this change makes the new version more general compared to its old version:

$$Employee \sqsubseteq Employee'$$

According to our procedure, this should not be a problem because employee is in the “subsuming list”.

When we analyze this change, we see that it has an impact on the definition of the concept “DepartmentMember” as it enlarges the set of objects allowed to take the first place in the has-member relation. This leads to a new definition of *DepartmentMember'* with $DepartmentMember \sqsubseteq DepartmentMember'$. As “DepartmentMember” was already more general than “HeadOfDepartment” and the employee concept is not used in the definition of the latter the, implied subsumption relation indeed still holds.

5.3.2 Example 2: the department-manager relation

The second example, we have to deal with a change affecting a relation that is used in an external definition. The relation department-manager is specialized by restricting its range to the concept “manager” (which is a subconcept of employee) making it a subrelation of its previous version:

$$\text{department} - \text{manager} \sqsupseteq \text{department} - \text{manager}'$$

Again, this is harmless according to our procedure, as department_manager is in the “subsumed list”.

The analysis shows that this change has an impact on the definition of the concept “HeadOfDepartment” as it restricts the allowed objects to the more specific concept “Manager”. The new definition *HeadOfDepartment'* is more specific than the old one: $\text{HeadOfDepartment}' \sqsubseteq \text{HeadOfDepartment}$. As the old version was already more specific than the definition of “DepartmentMember” and the “department-manager” relation is not used in the definition of the latter the implied subsumption is indeed still valid.

The situation is different if the range of the “department-manager” relation is changed to the concept “person” which is more general than “employee”. In this case the definition of the concept “department-manager” also becomes more general. This means that we cannot guarantee that it is still subsumed by “DepartmentMember”. In this case we have to recompute and compile the implied subsumption relations in order to guarantee integrity.

5.3.2.1 Example 3: the department concept The different changes of the definition of the “department” concept left us with no clear idea of the relation between the old and the new versions. In this specific case, however, we can still make assertions about the impact on implied subsumption relations. The reason is that the concept occurs in both definitions. Moreover, it plays the same role, namely restricting the domain of the relation that connects an organizational unit with the set of objects that make up the externally defined concept. As a consequence, the changes have the same impact on both definitions, thus not invalidating the implied subsumption relation.

6 Summary

In this article, we discussed an infrastructure for the representation of and reasoning with modular ontologies. The intention was to enhance the existing semantic web infrastructure with notions of modularization that have been proven useful in other areas of computer science, in particular in software engineering. We defined a set of requirements for modular ontologies that arise from expected benefits such as enhanced re-use and more efficient reasoning. Taking the requirements of loose coupling, self containment and integrity as a starting point, we defined a framework for modular ontologies providing the following contributions to the state of the art in ontology representation for the semantic web:

- (1) We presented a formal model for describing dependencies between different ontologies. We proposed conjunctive queries for defining concepts using elements from another ontology and presented a model-based semantics in the spirit of Distributed Description Logics that provides us with a notion of logical consequence across different ontologies.
- (2) We compared our model with the existing standard, i.e. the web ontology language OWL and showed that the OWL import facilities can easily be captured as a special case in our model. We further showed that our model provides additional expressiveness in particular with respect to modeling relations. In order to get a better idea of the improvements of our model over OWL, we investigated the formal properties of inter module mappings, their impact on reasoning and their intuition.
- (3) We described a method for detecting changes in an ontology and for assessing their impact. The main feature of this method is the derivation of conceptual changes from purely syntactic criteria. These conceptual changes in turn provide input for a semantical analysis of the effect on dependent ontologies, in particular on the validity of implied subsumption relations. We applied the method in a case study in the Wonder Web project and were able to determine the impact of changes without logical reasoning.

7 Discussion

There are three major questions connected to the approach for reasoning and managing change in modular ontologies proposed in this article. The first is *feasibility* in terms of computational complexity. As mentioned above, we use a heuristic approach to tackle this problem, which raises the question about the *adequacy* of the heuristics used. We argued that we chose a trade-off that works well in the context of OWL and typical semantic web applications.

This focus on a particular kind of representations finally raises the question of *generality* of the approach. We discuss these three basic questions in the following.

7.1 *Feasibility*

Reasoning in modular ontologies is complex. We have shown that the complexity is essentially the same as for reasoning in Classical Description Logics which are the basis for OWL. We cannot escape this complexity, but we can move parts of the reasoning effort offline by compiling implied subsumption relations as described in Section 3.2. This approach, however, is only feasible if there are phases where the offline computation necessary to compile the implied relations can be done without affecting the performance of the system. Typically, such computations are done ‘overnight’ when the system load can be assumed to be low. An alternative for situations where this approach is not possible is to do the compilation on the fly. In particular, we can compile implied subsumption relations whenever they are computed in order to answer a user query to the system. This kind of ‘lazy compilation’ has the advantage that the enormous effort for compiling implied knowledge is done as part of the normal reasoning process. In the beginning, users will not benefit much from this approach, but the time savings increase with each query answered. In this way, we also prevent compiling knowledge that is never used.

The main problem connected with the compilation approach, which is also a central aspect of this paper is the integrity of the compiled knowledge. In general compilation approaches only pay off if the computation time saved by being able to use compiled knowledge is not larger than the effort of updating the compiled knowledge. This means, that compilation only makes sense in rather stable systems. In principle, we can assume that knowledge on the terminological level as it is represented in ontologies is normally more stable than instance data as normally found in databases. While changes to ontologies will occur less frequently they can still have a significant impact on the system. For this reason, our work focussed on heuristics for efficiently updating the system when changes occur. In this context, the feasibility of the approach relies on the adequacy of the heuristics chosen.

7.2 *Adequacy*

Our method for detecting harmful changes is a conservative one. We basically identify changes that are obviously not harmful and refrain from updating when only such changes occurs. The criterion used for this purpose is a rather

weak one as we do not consider the specific changes made to a concept but restrict our analysis to the semantic relation between the old and the new version of the concept definition. The advantage of this choice is the efficiency of the approach as the criterion for harmless changes can be checked in linear time with respect to the number of concept names involved. We further weaken the method by not actually computing the semantic relation between the old and the new version of the concept, but rather determine the relation based on a set of change operations determined by syntactic analysis. This method also constitutes an incomplete heuristic as for some changes, we cannot determine the relation between the old and the new version. Again, the advantage of this choice lies in the efficiency of the approach, because it allows us to live completely without Description Logic reasoning.

The question that arises is whether the degree of incompleteness of our approach is justified. Currently there are no experimental results that support the usefulness of the heuristics, however, it is clear that the heuristics cover a wide range of relevant cases. On the level of determining the semantic relation between the versions, the library of change operations covers all possible change operations and most of them have a known effect that can be exploited in our approach. The cases in which the effect is not known can be assumed to be the hard cases that will always require complete reasoning independent of the heuristics used. For this reason, we believe that we cannot be much better on this level without falling back to classical subsumption reasoning. On the level of determining harmless change, the heuristics used is quite weak as well. In particular we claim that a change is only harmless if all concepts and relations satisfy certain properties. This can certainly be relaxed if we invest more time in the analysis of the actual definitions of the concepts involved. For example, changes in concepts that occur in both, the subsumed and the subsuming concepts are not relevant due to the monotonicity of the effect. We could also try to determine which parts of the definitions actually contribute to the subsumption proof and consider changes in other parts of the definitions as harmless as well. These improved heuristics still fit into the approach proposed and are subject of future work. For the time being, we conclude that the general mechanisms are in place and that the heuristic described in this paper already covers many relevant cases as we show in the examples from the WonderWeb case study. In principle, the choice of the best heuristic will rely on the specific application scenario and in particular on the actual expressiveness used in the models.

7.3 Generality

A final point for discussion is the generality of the approach described. Throughout this paper, we based our discussions on Description Logics as a representation language for ontologies, Distributed Description Logics for providing the semantics of mappings as well as the equivalent of conjunctive queries for describing relations between different modules. All of these choices are carefully made and are motivated by practical as well as theoretical considerations. Probably the most uncontroversial choice is the one for Description Logics for encoding ontologies. In the context of semantic web research, Description Logics have become the primary language for describing terminological knowledge mostly in terms of the Web Ontology Language OWL. Our approach covers most of the expressiveness of OWL-DL with the exception of nominals. As a result, most existing OWL ontologies will fit in our framework and could easily be turned into modular ontologies by adding external concepts.

A choice that is less obvious is Distributed Description Logics as a basis for the semantics of mappings. In a recent survey, we compared different approaches for describing mapping semantics [35]. One result of this comparison was that Distributed Description Logics provide the highest degree of de-coupling between different T-Boxes. This is important for our purposes as we want to support localized reasoning. A generalization of Distributed Description Logics in terms of a distributed version of first order logic has been described by Serafini and Ghididi [36]. We could have chosen this more general framework as the basis for our work, however, the drawback of this is the lack of existing reasoning methods. Distributed Description Logics come with a well investigated and implemented proof system that can be used to implement our approach.

The most discussable choice is to restrict the language that can be used to define external concepts. The framework of Distributed Description Logic allows us to use arbitrary *SHIQ* expressions in the definitions. A corresponding more general approach would have the same properties with respect to logical consequence, compilation and local reasoning. The restriction to the equivalent of conjunctive queries was motivated by the importance of the monotonicity property for the definition of update heuristics. This means that external concepts can be defined using a more expressive language. This, however, would come at the price that implied subsumption relations concerning this concept would have to be recomputed every time a change occurs. We believe that the restriction proposed in this paper is reasonable as it allows the use of update heuristics and also resembles view-based information integration, which is the dominant approach for describing mappings between database schemata, which also use conjunctive queries for describing mappings.

References

- [1] S. Bechofer, I. Horrocks, C. Goble, R. Stevens, OilEd: A reason-able ontology editor for the semantic web, in: F. Baader, G. Brewka, T. Eiter (Eds.), KI 2001: Advances in Artificial Intelligence, Springer, 2001, pp. 396–408.
- [2] N. Noy, R. Fergerson, M. Musen, The knowledge model of protege-2000: Combining interoperability and flexibility (2000).
URL citeseer.nj.nec.com/noy01knowledge.html
- [3] I. Horrocks, The FaCT system, in: H. de Swart (Ed.), Automated Reasoning with Analytic Tableaux and Related Methods: International Conference Tableaux'98, no. 1397 in Lecture Notes in Artificial Intelligence, Springer-Verlag, Berlin, 1998, pp. 307–312.
URL download/1998/t98-paper.ps.gz
- [4] V. Haarslev, R. Moller, Description of the RACER system and its applications, in: Proceedings of the Description Logics Workshop DL-2001, Stanford, CA, 2001, pp. 132–142.
- [5] J. Broekstra, A. Kampman, F. van Harmelen, Sesame: A generic architecture for storing and querying rdf and rdf schema, in: The Semantic Web - ISWC 2002, Vol. 2342 of Lecture Notes in Computer Science, Springer, 2002, pp. 54–68.
- [6] W. Farmer, J. Guttman, F. Thayer, Little theories, in: D. Kapur (Ed.), Proceedings of the Eleventh International Conference on Automated Deduction, Vol. 607 of Lecture Notes in Computer Science, Springer Verlag, 1992, pp. 567–581.
- [7] P. Clark, J. Thompson, K. Barker, B. Porter, V. Chaudhri, A. Rodriguez, J. Thomere, S. Mishra, Y. Gil, P. Hayes, T. Reichherzer, Knowledge entry as the graphical assembly of components, in: Proceedings of the 1st International Conference on Knowledge Capture (K-Cap'01), 2001.
- [8] E. Amir, S. McIlraith, Partition-based logical reasoning, in: 7th International Conference on Principles of Knowledge Representation and Reasoning (KR'2000), 2000.
- [9] S. McIlraith, E. Amir, Theorem proving with structured theories, in: B. Nebel (Ed.), Proceedings of IJCAI'01, Morgan Kaufmann, San Mateo, 2001, pp. 624–634.
- [10] S. Lauritzen, D. Spiegelhalter, Local computations with probabilities on graphical structures and their application to expert systems, *Journal of the Royal Statistical Society* 50.
- [11] A. Rector, Modularisation of domain ontologies implemented in description logics and related formalisms including OWL, in: Proceedings of the 16th International FLAIRS Conference, AAAI, 2003.

- [12] M. Buchheit, F. D. W. Nutt, A. Schaerf, Terminological systems revisited: Terminology = schema + views, in: Proceedings of the 12th National Conference on Artificial Intelligence (AAAI-94), 1994.
- [13] F. Giunchiglia, C. Ghidini, Local models semantics, or contextual reasoning = locality + compatibility, in: Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98), Morgan Kaufmann, 1998, pp. 282–289.
- [14] A. Borgida, L. Serafini, Distributed description logics: Directed domain correspondences in federated information sources, in: R. Meersman, Z. Tari (Eds.), On The Move to Meaningful Internet Systems 2002: CoopIS, Doa, and ODBase, Vol. 2519 of LNCS, Springer Verlag, 2002, pp. 36–53.
- [15] D. Brickley, R. Guha, A. Layman, Resource description framework (RDF) schema specification, Working draft, W3C, <http://www.w3c.org/TR/WD-rdf-schema> (August 1998).
- [16] M. Dean, D. Connolly, F. van Harmelen, J. Hendler, I. Horrocks, D. McGuinness, P. Patel-Schneider, L. Stein, Web ontology language (owl) reference version 1.0, Working draft, W3C, <http://www.w3.org/TR/owl-ref/> (November 2002).
- [17] R. Volz, A. Mdche, D. Oberle, Towards a modularized semantic web, in: Proceedings of the ECAI'02 Workshop on Ontologies and Semantic Interoperability, 2002.
- [18] R. Volz, D. Oberle, R. Studer, Views for light-weight web ontologies, in: Proceedings of the ACM Symposium on Applied Computing SAC 2003, 2003.
- [19] I. Horrocks, U. Sattler, S. Tobies, Reasoning with individuals for the description logic *SHIQ*, in: Proceedings of the 17th International Conference on Automated Deduction (CADE-17), Lecture Notes in Computer Science, Springer Verlag, 2000.
- [20] A. Borgida, L. Serafini, Distributed description logics: Assimilating information from peer sources., *Journal of Data Semantics* 1 (2003) 153–184.
- [21] L. Serafini, A. Borgida, A. Tamilin., Aspects of distributed and modular ontology reasoning., in: Proceedings of the International Joint Conference on Artificial Intelligence - IJCAI-05, Edinburgh, Scotland, 2005.
- [22] F. Baader, D. Calvanese, D. McGuinness, D. Nardi, P. Patel-Schneider (Eds.), *The Description Logic Handbook - Theory, Implementation and Applications*, Cambridge University Press, 2003.
- [23] J. D. Ullman, Information integration using logical views, *Theoretical Computer Science* 239 (2) (2000) 189–210.
URL citeseer.ist.psu.edu/ullman97information.html
- [24] R. Fikes, P. Hayes, I. Horrocks, Owl-ql: A language for deductive query answering on the semantic web, *Journal of Web Semantics* 2 (1).

- [25] I. Horrocks, S. Tessaris, A conjunctive query language for description logic aboxes, in: AAI/IAAI, 2000, pp. 399–404.
URL citeseer.nj.nec.com/horrocks00conjunctive.html
- [26] S. Tobies, Complexity results and practical algorithms for logics in knowledge representation., Phd thesis, LuFG Theoretical Computer Science, RWTH-Aachen (2001).
- [27] L. Serafini, A. Tamilin, DRAGO: Distributed reasoning architecture for the semantic web, in: In Proceedings of the Second European Semantic Web Conference (ESWC'05), Springer-Verlag, 2005.
- [28] J. Banerjee, W. Kim, H.-J. Kim, H. F. Korth, Semantics and Implementation of Schema Evolution in Object-Oriented Databases, SIGMOD Record (Proc. Conf. on Management of Data) 16 (3) (1987) 311–322.
URL <http://doi.acm.org/10.1145/38713.38748>
- [29] M. Klein, Change management for distributed ontologies, Ph.D. thesis, Vrije Universiteit Amsterdam (Aug. 2004).
URL <http://www.cs.vu.nl/~mcaklein/thesis/>
- [30] E. Franconi, F. Grandi, F. Mandreoli, A semantic approach to schema evolution and versioning in object-oriented databases, in: Proceedings of CL 2000, Vol. 1861 of Lecture Notes in Artificial Intelligence, Springer Verlag, 2000, pp. 1048–1062.
- [31] M. Klein, D. Fensel, A. Kiryakov, D. Ognyanov, Ontology versioning and change detection on the web, in: 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW02), no. 2473 in LNCS, Sigüenza, Spain, 2002, p. 197 ff.
URL <http://www.cs.vu.nl/~mcaklein/papers/EKAW02.pdf>
- [32] A. Gangemi, N. Guarino, C. Masolo, A. Oltramari, L. Schneider, Sweetening ontologies with DOLCE, in: 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW02), Vol. 2473 of Lecture Notes in Computer Science, Sigüenza, Spain, 2002, p. 166 ff.
URL <http://link.springer.de/link/service/series/0558/bibs/2473/24730166.htm>
- [33] C. Masolo, S. Borgo, A. Gangemi, N. Guarino, A. Oltramari, Ontology library (final), Deliverable D18, EU/IST Project WonderWeb (Dec. 2003).
URL <http://wonderweb.semanticweb.org/deliverables/documents/D18.pdf>
- [34] R. Volz, D. Oberle, S. Staab, R. Studer, Ontolift prototype, Deliverable D11, EU/IST Project WonderWeb (2002).
- [35] L. Serafini, H. Stuckenschmidt, H. Wache, A formal investigation of mapping languages for terminological knowledge, in: Proceedings of the 19th International Joint Conference on Artificial Intelligence - IJCAI05, Edingurgh, UK, 2005.
- [36] C. Ghidini, L. Serafini, Distributed first order logic - revised semantics, Technical report, ITC-irst (January 2005).