



UNIVERSITY  
OF TRENTO

---

DEPARTMENT OF INFORMATION AND COMMUNICATION TECHNOLOGY

---

38050 Povo – Trento (Italy), Via Sommarive 14  
<http://www.dit.unitn.it>

QUERY PROCESSING IN ONTOLOGY-BASED  
PEER-TO-PEER SYSTEMS

Heiner Stuckenschmidt, Frank van Harmelen  
and Fausto Giunchiglia

November 2002

Technical Report # DIT-02-0096



# Query Processing in Ontology-Based Peer-to-Peer Systems

Heiner Stuckenschmidt, Frank van Harmelen  
Artificial Intelligence Department, Vrije Universiteit Amsterdam  
Fausto Giunchiglia

Department of Information and Communication Technology,  
University of Trento

November 28, 2002

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Semantic Web and Peer-to-Peer . . . . .	1
1.2	Related Work . . . . .	2
1.3	The Need for New Approaches . . . . .	2
<b>2</b>	<b>Ontology-Based Peer-to-peer Systems</b>	<b>5</b>
2.1	Ontological Knowledge . . . . .	6
2.2	Inter-Ontology Mappings . . . . .	8
2.3	Semantics and Logical Consequence . . . . .	9
2.4	Ontology-Based Queries . . . . .	10
<b>3</b>	<b>Query Processing</b>	<b>12</b>
3.1	Approximating Concepts . . . . .	12
3.2	Queries as Concepts . . . . .	14
3.3	Quality of Approximation . . . . .	16
<b>4</b>	<b>Query Relaxation</b>	<b>18</b>
4.1	Variable Elimination . . . . .	19
4.2	Guided Elimination . . . . .	19
<b>5</b>	<b>Examples from a Case Study</b>	<b>22</b>
5.1	Concept Approximations . . . . .	22
5.2	Query Relaxation . . . . .	24
<b>6</b>	<b>Conclusions</b>	<b>26</b>

## Abstract

The distributed, heterogeneous and unstructured nature of the Web poses a new challenge to query-answering over multiple data sources. In particular, it is no longer realistic to assume that the involved data sources act as if they were a single (virtual) source, modeled as a global schema, as is done in classical data integration approaches. In this paper, we propose an alternative approach where we replace the role of a single virtual data source schema with a peer-to-peer approach relying on limited shared (or: overlapping) vocabularies between peers. Since overlaps between vocabularies of peers will be limited, query processing will have to be approximate. We provide a formal model for such approximate query processing based on limited shared vocabularies between peers, and we show how the quality of the approximation can be adjusted in a gradual manner. The result is a flexible architecture for query-processing in large, distributed and heterogeneous environments, based on a formal foundation. This architecture is suitable for knowledge-sharing in the peer-to-peer-style networks that are expected to be typical of the Semantic Web.

**Keywords:** Semantic Web, Methods and Formalisms for Knowledge Sharing, Knowledge-Based Mediation Architectures

# Chapter 1

## Introduction

### 1.1 Semantic Web and Peer-to-Peer

The approach to query-processing that we present in this paper is strongly motivated by the peer-to-peer (P2P) architecture [Parameswaran et al., 2001] that we expect for the Semantic Web. In this section we will argue why we expect the Semantic Web to have such a peer-to-peer architecture.

When we look at the current World Wide Web, we see in fact a mixed architecture, that is partly client/server-based, and partly P2P. On the one hand, each node in the network can directly address every other node in the network in a single, flat, world-wide address space, giving it the structure typical of many P2P networks. On the other hand, in practice there is currently a strong asymmetry between nodes in this address space that acts as content-servers, and nodes that act as clients. Recent estimates indicate the presence of 50 million web-servers, but as many as 150 million clients. On the scale of the World Wide Web, any form of centralization would create immediate bottle necks, in terms of network throughput and server capacity.

This need for a flat, non-server-centered architecture will be even stronger on the Semantic Web. Of course, the same physical load-balancing arguments hold as on the current Web, but the Semantic Web adds a new argument in favor of a P2P-style argument. On the Semantic Web, any server-centered architecture will not only create physical bottlenecks, but will also create *semantic bottle-necks*. Since the semantics of information will be explicit (or at least: more explicit) on the Semantic Web, any single server will in a way “impose” a particular semantic view on all its clients. This will have undesirable consequences, both in terms of the pluriformity of the available information, as well as in terms of the size of the ontology that such information-servers would have to maintain.

Instead, a P2P-style architecture will be able to avoid both the physical

and the semantic bottle-necks. Different semantic views, expressed in terms of different ontologies, will be provided by many peers in a flat network of peers, each employing their own local, small ontology.

In the next subsection ??, we will briefly survey the existing work on integration of heterogeneous databases, and we will see that this work is predominantly based on the notion of a global schema that is connected to the heterogeneous schemas to be integrated. Subsequently, in section ?? we will argue why this traditional approach is no longer viable in a peer-to-peer style network as the Semantic Web will be. The remainder of the paper will then be devoted to describing our proposal for such new approaches that will enable us to do query processing in a peer-to-peer setting without the need for global integrating schemas.

## 1.2 Related Work

The problem of integrating heterogeneous database schemas has been addressed by many researchers (see [Levy, 1999] or for a survey). The integration is normally done using a global schema that is connected to the heterogeneous schemas to be integrated by a number of views. We can distinguish two general approaches:

- *Global-As-View*: In the global-as-view approach every relation in the global schema is defined as a view over the different schemas to be integrated (see e.g. [Garcia-Molina et al., 1997]).
- *Local-As-View*: In the local-as-view approach, views are used to define the schemas of local information sources in terms of the global one [Levy et al., 1996]).

The benefits of using explicit semantic models. i.e. ontologies has been recognized in many approaches. A survey of approaches using ontologies is provided by [Wache et al., 2001]. Description logics have been proven to be a useful formalism for specifying and reasoning about semantic models [Calvanese et al., 2001] to support information integration. It has been shown that results from the database area provide solutions for the integration semi-structured information, i.e. of XML documents, (see e.g. [Bergamaschi et al., 1999]). The idea of rewriting representations based on the special capabilities of a remote system is reported in [Papakonstantinou et al., 1996] and has been applied for translating SQL queries [Chang and Garcia-Molina, 2001]. The use of description logics for this rewriting is described in [Goasdoué and Rousset, 2000].

### 1.3 The Need for New Approaches

Consider the situation where John, a person living in Toronto, is described in the database F of his family doctor, and also in the database H of the hospital where he once received medical treatment.

Example: John goes to a ski resort in another country, for instance Trentino in Italy. Unluckily, here he has an accident; for instance, he breaks a leg, and he must get medical aid to the resort's medical office. This office, has its own database M which now needs to get involved. M may need to query H for the purpose of retrieving treatment details of a similar past accident. Furthermore, when John returns home, a new record from M should appear in F. However the acquaintance between M and F does not need to be maintained for ever, since the two databases will probably not need to coordinate again, and can eventually be dropped.

#### **Dropping the global schema**

In situations like that described in this example, the design and development of data integration mechanisms for randomly acquainted databases which may need to communicate only a few times, becomes impractical.

In this context, it makes very little sense to speak of a global schema, as it is commonly done in the data integration literature [Halevy, 2001]. The main conceptual reason for rejecting this notion is that we cannot think of a set of P2P databases just as a particular implementation of a single (virtual) database, this being the underlying assumption which motivates the definition of a global schema.

From a foundational point of view, any theory developed under the assumption of a global schema, and under the implicit assumption that the global schema is fixed, prevents us from the studying the dynamics of a P2P network. As far as we know these two assumptions have never been relieved, in particular in the data integration literature, see for instance [Halevy, 2001, Lenzerini, 2002]. It is no longer possible to see the global schema as a view of the local database (global-as-view approach) or, vice versa, the local databases as views of the global database (local-as-view approach). For instance, we can no longer assume that there is a unique universe containing all the elements of the single databases, but rather many overlapping domains.

#### **Good enough answers**

In the setting described above, it becomes hard to maintain a high quality level in the answers that the P2P network is able to provide. By high quality level we mean the fact that data can flow among the databases preserving (at the best possible level of approximation) soundness and completeness. In this



context, soundness means that the data provided by the local databases satisfy the global schema (but they are not necessarily complete, some of them can get lost in the coordination). Completeness has the dual meaning. In the data integration literature, completeness is often given up, still maintaining the request of soundness. In a P2P environment, completeness and soundness will be very hard to achieve. This will happen in limit cases, for instance with low dynamics, simplified interaction among the databases, and if and when there will be interest in investing a substantial amount of money in the solution of the problem.

One area where there will often be interest in getting very high quality data integration is the medical care domain. There are however many other application domains where this is not the case. One such example is tourism. This domain, is not life critical, and in many cases the small dimension of a single business (hotel, campsite, ?) does not justify big investments. Consider the following example.

Example: When planning his vacation in Trentino, John goes to a local agency, which unluckily can not offer John anything from their own database. Instead the agency searches for single operators in the Trentino region (hotels, ski resorts, etc), starts communication sessions with some of them, and queries for the necessary information (e.g., prices, conditions, availability).

Compared with the medical care example, the dynamics will have a much higher impact on the quality of the answer. We have network variance: the relevant databases are much more unstable in their being active and coordinated in the network, nodes come and go (for instance depending on the season), and so on. We have database variance: John travels around and queries different databases. The same query will get different results since each database will implement different degrees of coordination with the others, and so on. Thus, for instance, a query about hotels made to a hotel database will likely get an answer that is better than the answer obtained from a campsite database. We also have query variance: if you ask a query about campsites to a campsite database you will likely get a better quality answer than if you ask this database a query about hotels. Depending on the query, certain coordination mechanisms may or may not be activated. However, in this application, the agency doesn't need the best possible answer. It simply needs some answer. As long as, for instance, it gets a hotel John likes, this is good enough. Compared to the previous example, much lower quality data coordination will probably suffice.

The medical care and tourism domains are just examples. Things can get even more radical and complex when one thinks of applications where some of the nodes are mobile and where coordination happens on an even more occasional basis, for instance due to the physical proximity of two mobile peers. In these situations, and for certain kinds of applications, almost any answer will suffice, as long as there is one. In a P2P environment, in terms of quality of

answers, it is possible to go from one extreme to the other. On one extreme, it may be usual to get poor quality answers. This may happen because the databases interact partially or do not interact at all or, even worse, they pass around data which are wrong (for instance because of unsolved problems of semantic heterogeneity). On the other extreme, there will be a tight coordination and it will be possible to achieve or, at least, approximate soundness and completeness. Between these two extremes there is a continuum of answers of different quality.

## Chapter 2

# Ontology-Based Peer-to-peer Systems

Before we can present our approach to query processing in ontology-based Peer-to-Peer systems, we have to specify the systems we are talking about in more detail. We assume a system of independent peers that encapsulate the (possibly redundant) information of the whole system. Each peer uses one or model ontologies. These ontologies are used as a conceptual schema of the actual information that can be seen as an instantiation of the ontology. Peers exchange knowledge by formulating queries using the vocabulary defined in the ontologies they use and sending them the other peers in the network. The task of the receiving peer is to determine the answer to this queries relative to its own vocabulary and information. This leads to a situation where we are rather concerned with heterogeneous knowledge bases that plain data sources in the conventional sense.

In order to get a clearer notion of the problem of processing queries in such systems we make some simplifying assumptions. First of all we will only consider two peers that want to communicate. Then we assume that there are only two ontologies involved, a shared one and a private one of the peer trying to communicate. We further assume that both ontologies are encoded on the same language, preventing us from the problem of integrating the ontology languages. Figure 2.1 illustrated the situation.

This simplified communication problem can easily be extended to more realistic scenarios as communication is mostly bi-lateral even in complex multi-agent systems. There might be more than two ontologies involved in the communication, but they will all either be shared or private to one on the agents. The assumption that there are actually ontologies being shared by peers in the system is backed by the observation, that real-world ontologies are in most cases not build from scratch. It is rather common to at least start with

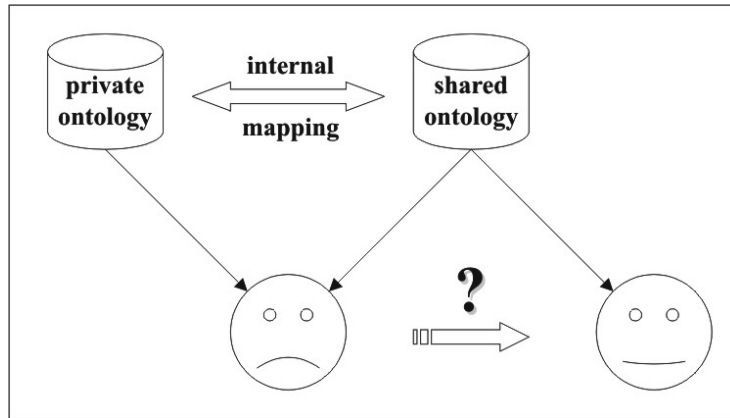


Figure 2.1: The communication problem

an existing ontology (see for example <http://www.daml.org/> for a library of ontologies about various domains. The existence of an internal mapping between the ontologies used by an individual agent is likely because if the peers wants to use more than one ontology as a basis for its information, it has to know about their relation. As a single peer is a rather static system compared to the overall network, we can use schema matching techniques that have been developed in the database community in order to find correspondences (see [Rahm and Bernstein, 2001] for an overview). The only assumption that really is a simplification is the existence of a single ontology language. Investigating this problem, however, is out of the scope of this paper. For an approach to overcome language heterogeneity, we refer to [Euzenat, 2001].

In the following we give formal definitions for the parts of an ontology-based peer-to-peer system that are concerned with query processing. These parts include the definition of ontologies and mappings as well as the notion of queries and answers in the setting of ontology-based information.

## 2.1 Ontological Knowledge

A number of languages for encoding ontologies on the Web have been proposed (see [Gomez-Perez and Corcho, 2002] for an overview). In order to get a general notion of ontological knowledge, we define the general structure of a terminological knowledge base (ontology) and its instantiation independent of a concrete language.

**Definition 1 (Terminological Knowledge Base)** *A Terminological Knowledge Base  $T$  is a triple*

$$\mathcal{T} = \langle \mathcal{C}, \mathcal{R}, \mathcal{O} \rangle$$

where  $\mathcal{C}$  is a set of class definitions,  $\mathcal{R}$  is a set of relation definitions and  $\mathcal{O}$  is a set of object definitions.

Terminological knowledge usually groups objects of the world that have certain properties in common (e.g. cities or countries). A description of the shared properties is called a class definition. Concepts can be arranged into a subclass-superclass relation in order to be able to further discriminate objects into subgroups (e.g. capitals or European countries). Classes can be defined in two ways, by enumeration of its members or by stating that it is a refinement of a complex logical expressions. The specific logical operators to express such logical definitions can vary between ontology languages; the general definitions we give here abstract from these specific operators.

**Definition 2 (Class Definitions)** *A class definition is an axiom of one of the following forms:*

- $c \equiv (o_1, \dots, o_n)$  where  $c$  is a class definition and  $o_1, \dots, o_n$  are object definitions.
- $c_1 \sqsubseteq c_2$  where  $c_1$  and  $c_2$  are class definitions.

Further, there is the universal class denoted as  $\top$ .

Objects of the same type normally occur in similar situations where they have a certain relation to each other (cities lie in countries, countries have a capital). These typical relations can often be specified in order to establish structures between classes. Terminological knowledge considers binary relations that can either be defined by restricting their domain and range or by declaring it to be a sub-relation of an existing one.

**Definition 3 (Relation Definitions)** *A relation definition is an axiom of one of the following forms:*

- $r \sqsubseteq (c_1, c_2)$  where  $r$  is a role definition and  $c_1$  and  $c_2$  are class definitions.
- $r_1 \sqsubseteq r_2$  where  $r_1$  and  $r_2$  are role definitions.

The universal role is defined as  $\top \times \top$ .

Sometimes single objects (e.g. the continent Europe) play a prominent role in a domain of interest, or the membership of a concept is defined by the relation to a specific object (European countries are those contained in Europe). For this purpose ontology languages often allow to specify single objects, also called instances. In our view on terminological knowledge, instances can be defined by stating their membership in a class. Further, we can define instances of binary relations by stating that two objects form such a pair.

**Definition 4 (Object Definitions)** *An object definition is an axiom of one of the following forms:*

- $o : c$  where  $c$  is a class definition and  $o$  is an individual.
- $(o_1, o_2) : r$  where  $r$  is a relation definition and  $o_1, o_2$  are object definitions.

In the following, we will consider terminological knowledge bases that consist of such axioms. Of course, any specific ontology language will have to further instantiate these definitions to specify logical operators between classes etc, but for the purposes of this paper, these general definitions are sufficient. Further, we define the signature of a terminological knowledge base to be a triple  $\langle \mathcal{CN}, \mathcal{RN}, \mathcal{IN} \rangle$ , where  $\mathcal{CN}$  is the set of all names of classes defined in  $\mathcal{C}$ ,  $\mathcal{RN}$  the set of all relation names and  $\mathcal{IN}$  the set of all object names occurring in the knowledge base.

## 2.2 Inter-Ontology Mappings

We assume that each peer has an integrated view on the ontologies it uses as a semantic foundation for its information. This integrated view is created by the mappings relate elements from different ontologies. The creation of these mappings is discussed in other work (see e.g. [Rahm and Bernstein, 2001] for an overview) and is not further discussed in this paper. As our methods for query processing rely on these internal mappings of individual peers, we have to define the nature of mappings. For this work, we adopted the mapping framework proposed in [Madhavan et al., 2002] summarized in the following.

Madhavan et.al. define mappings in terms of operations between expressions in two different domain models, in our case ontologies. They further demand that the resulting expression is consistent with the logical interpretation of the individual ontologies (we cover this point in the next section). As this framework is very generic, we instantiate it in the context of ontological knowledge as defined in the last section. In especially, we need to define the kind of expressions, we consider and the operations used to relate them.

The expressions, we are considering here are the definitions of concepts, relations and objects, respectively. Further, we define the following mappings between the different types of knowledge:

$$\begin{aligned}
 c_1 & \xleftrightarrow{m_C} c_2, m_C \in \{\sqsubseteq, \sqsupseteq, \equiv\} \\
 r_1 & \xleftrightarrow{m_R} r_2, m_R \in \{\sqsubseteq, \sqsupseteq, \equiv\} \\
 o_1 & \xleftrightarrow{\equiv} o_2
 \end{aligned} \tag{2.1}$$

Intuitively the mappings marked with the operator  $\equiv$  state that definitions in the different ontologies refer to the same concept, relation or individual. Mappings marked with  $\sqsubseteq$  (and  $\sqsupseteq$ ) state that the definition in the one ontology is a special

(more general) case of the other definition. Further, Madhavan et.al. consider the use of a helper model individual models are mapped to. The helper model is then used to derive composed mappings between these models. In our view on ontology-based peer-to-peer systems, the shared ontologies can be seen as such helper models. The formal definition of the mappings will be given in the next section.

## 2.3 Semantics and Logical Consequence

We can define semantics and logical consequence of a terminological knowledge base using an interpretation mapping  $\mathfrak{S}$  into an abstract domain  $\Delta$  such that:

- $c^{\mathfrak{S}} \subseteq \Delta$  for all class definitions  $c$  in the way defined above
- $r^{\mathfrak{S}} \subseteq \Delta \times \Delta$  for all relation definition  $r$
- $o^{\mathfrak{S}} \in \Delta$  for all object definitions  $o$

This type of denotational semantics is inspired by description logics [Donini et al., 1996], however, we are not specific about operators that can be used to build class definitions which are of central interest of these logics. Using the interpretation mapping, we can define the notion of a model in the following way:

**Definition 5 (Model of a Terminological Knowledge Base)** *An interpretation  $\mathfrak{S}$  is a model for the knowledge base  $\mathcal{T}$  if  $\mathfrak{S} \models A$  for every axiom  $A \in (\mathcal{C} \cup \mathcal{R} \cup \mathcal{O})$  where  $\models$  is defined as follows.*

- $\mathfrak{S} \models c \equiv (o_1, \dots, o_n)$ , iff  $c^{\mathfrak{S}} = \{o_1^{\mathfrak{S}}, \dots, o_n^{\mathfrak{S}}\}$
- $\mathfrak{S} \models c_1 \sqsubseteq c_2$ , iff  $c_1^{\mathfrak{S}} \subseteq c_2^{\mathfrak{S}}$
- $\mathfrak{S} \models r \sqsubseteq (c_1, c_2)$ , iff  $r^{\mathfrak{S}} \subseteq c_1^{\mathfrak{S}} \times c_2^{\mathfrak{S}}$
- $\mathfrak{S} \models r_1 \sqsubseteq r_2$ , iff  $r_1^{\mathfrak{S}} \subseteq r_2^{\mathfrak{S}}$
- $\mathfrak{S} \models o : c$ , iff  $o^{\mathfrak{S}} \in c^{\mathfrak{S}}$
- $\mathfrak{S} \models (o_1, o_2) : r$ , iff  $(o_1^{\mathfrak{S}}, o_2^{\mathfrak{S}}) \in r^{\mathfrak{S}}$

In order to be able to handle multiple ontologies, we have to define the interpretation mapping over different models and mappings between them. In the following, we only consider an interpretation for two ontologies and mappings between them. The definitions, however, can easily be extended to more than two ontologies.

First of all, we divide the interpretation mapping  $\mathfrak{S}$  into two sub-mappings  $\mathfrak{S}^1$  and  $\mathfrak{S}^2$  each defining the interpretation for one of the two ontologies in the way described above. Further, we define the interpretation of mappings between the ontologies in the following way:

$$\mathfrak{S} \models (c_1 \xrightarrow{m_C} c_2) \text{ iff } \begin{cases} c_1^{\mathfrak{S}^1} = c_2^{\mathfrak{S}^2} \text{ for } m_C = \equiv \\ c_1^{\mathfrak{S}^1} \subseteq c_2^{\mathfrak{S}^2} \text{ for } m_C = \sqsubseteq \\ c_1^{\mathfrak{S}^1} \supseteq c_2^{\mathfrak{S}^2} \text{ for } m_C = \sqsupseteq \end{cases} \quad (2.2)$$

$$\mathfrak{S} \models (r_1 \xrightarrow{m_R} r_2) \text{ iff } \begin{cases} r_1^{\mathfrak{S}^1} = r_2^{\mathfrak{S}^2} \text{ for } m_R = \equiv \\ r_1^{\mathfrak{S}^1} \subseteq r_2^{\mathfrak{S}^2} \text{ for } m_R = \sqsubseteq \\ r_1^{\mathfrak{S}^1} \supseteq r_2^{\mathfrak{S}^2} \text{ for } m_R = \sqsupseteq \end{cases} \quad (2.3)$$

$$\mathfrak{S} \models (c_1 \xrightarrow{\equiv} c_2) \text{ iff } o_1^{\mathfrak{S}^1} = o_2^{\mathfrak{S}^2} \quad (2.4)$$

These definitions enable us to perform reasoning across different ontologies using the notion of logical consequence:

**Definition 6 (Logical Consequence)** *An axiom  $A$  logically follows from a set of axioms  $\mathcal{S}$  if  $\mathfrak{S} \models \mathcal{S}$  implies  $\mathfrak{S} \models A$  for every model  $\mathfrak{S}$ . We denote this fact by  $\mathcal{S} \models A$ .*

In order to perform logical reasoning, we can use existing reasoning systems that have been build for reasoning about description logic knowledge bases. The system that has been used in the case study is the FaCT system [Horrocks, 1998].

## 2.4 Ontology-Based Queries

As mentioned in the beginning of this section, the peers in a system will exchange information by querying information from other peers in the network using terms from their own ontology. In the following we first define ontology based queries as well as the notion of answers to and relations between queries. We formalize queries in the following way: conjuncts of a query are predicates that correspond to classes and relations of the ontology. Further, variables in a query may only be instantiated by constants that correspond to objects in that ontology.

**Definition 7 (Terminological Queries)** *Let  $V$  be a set of variables disjoint from  $\mathcal{IN}$  then an terminological query  $Q$  over a knowledge base  $\mathcal{T}$  is an expressions of the form*

$$Q \leftarrow q_{1_i} \wedge \cdots \wedge q_{m_i}$$

where  $q_i$  are query terms of the form  $x : c$  or  $(x, y) : r$  such that  $x, y \in V \cup \mathcal{IN}$ ,  $C \in \mathcal{CN}$  and  $R \in \mathcal{RN}$ .

The following expression in an example query based on an ontology used in the case study described later. It asks for hotels in castles that are located in towns in Meclenburg and have less than 25 rooms. :



$$\begin{aligned}
Q(X) \leftarrow & X : Hotel \wedge \\
& (X, Z) : hat - Zimmer \wedge Z \leq 25 \wedge \\
& (X, W) : ist - in - Schloss \wedge W = ja \\
& (X, Y) : liegt - in - Ort \wedge \\
& (Y, V) : liegt - in - Land \wedge V = mecklenburg \wedge
\end{aligned} \tag{2.5}$$

The fact that all conjuncts relate to elements of the ontology allows us to determine the answer to terminological queries in terms of instantiations of the query that are logical consequences of the knowledge base it refers to:

**Definition 8 (Query Answers (Halevy 2001))** *The answer of a query  $Q$  containing variables  $v_1, \dots, v_k$  over a knowledge base  $T$  is a set of tuples  $(i_1, \dots, i_k)$  such that  $T \models Q'$  where  $Q'$  is the query obtained from  $Q$  by substituting  $v_1, \dots, v_k$  by  $i_1, \dots, i_k$ . The projections of the answer tuples to variables occurring in the head of a query  $Q$  is denoted as  $res(Q)$  and referred to as the answer set of  $Q$ .*

The computation of query answers in the sense being defined above is the main inference task of peers within a system. In the next section we will discuss a logically well-founded approach for computing such answers.

## Chapter 3

# Query Processing

After having introduced some basic notions related to ontology-based peer-to-peer systems we now turn our attention to the problem of query processing. The definition of query answers provides us with a deductive definition that describe correct answers with respect to a query over a single specific ontology. In the case of an ontology-based Peer-to-peer system, however, we face a situation where we have to deal with more than one ontology. In many cases, the answering peer does not know all terms used in the query expression, because they are taken from the local ontology of the asking peer. In order to overcome this problem, we have to align the vocabularies the asking and the answering peer. In [Stuckenschmidt, 2002a] we describe an approach for approximately translating concept expressions from one ontology into another. We briefly summarize this approach in the next section and extend it to conjunctive queries. Further, we discuss the issue of the quality of approximations generated by our approach.

### 3.1 Approximating Concepts

The notion of an interpretation given above is a very general one and does not restrict the set of objects in the extension of a concept. This is done by the use of operators for defining classes. These kinds of operators restrict the possible extensions of a concept. Figure 3.1 defines some operators we use in the following in order to define classes.

Operator	Extension <sup>3</sup>
$C_1 \sqcap C_2$	$C_1^3 \cap \dots \cap C_n^3$
$\{x_1, \dots, x_n\}$	$\{x_1, \dots, x_n\} \subset \Delta$
$(\exists P.C)$	$\{y \in \Delta \mid \exists x ((y, x) \in P^3) \wedge x \in C^3\}$
$(\leq n)$	$\{x^3 \in N \mid x^E \leq n\}$

Figure 3.1: Some operators for Constraining Classes

These kinds of restriction are the basis for deciding whether a class definition is equivalent, more specialized or more general than another. Formally, we can decide whether one of the following relations between two expressions can be deduced from the ontology:

**subsumption:**  $C_1 \sqsubseteq C_2 \iff C_1^{\mathfrak{S}} \subseteq C_2^{\mathfrak{S}}$

**membership:**  $x : C \iff x^{\mathfrak{S}} \in C^{\mathfrak{S}}$

The classes in an ontology form a hierarchy with respect to the subsumption relation. In the case of multiple ontologies connected by mappings such a hierarchy can also be computed for the complete set of concepts. Therefore, we will always have a set of direct super- and a set of direct subclasses of a class  $c_1$  from the private ontology. We can use those direct sub- and super classes that belong to the shared ontology as upper and lower approximation for  $c_1$  in the shared ontology:

**Definition 9 (Lower Approximation)** *Let  $C_1$  be a set of private concepts,  $C_2$  a set of shared concepts of an agent and  $c \in C_1$  a class, then a class  $c_{glb} \in C_2$  is called a lower approximation of  $c$  in  $IS_2$ , if the following assertions hold:*

1.  $c_{glb} \sqsubseteq c$
2.  $(\exists c' \in C_2 : c' \sqsubseteq c) \implies (c' \sqsubseteq c_{glb})$

The greatest lower bound  $glb_{IS_2}(c)$  denotes the set of all lower approximations of  $c$  in  $C_2$ .

**Definition 10 (Upper Approximation)** *Let  $C_1$  be a private classes,  $C_2$  a set of shared classes of an agent and  $c \in C_1$  a private class, then a class  $c_{lub} \in C_2$  is called an upper approximation of  $c$  in  $IS_2$ , if the following assertions hold:*

1.  $c \sqsubseteq c_{lub}$
2.  $(\exists c' \in C_2 : c \sqsubseteq c') \implies (c_{lub} \sqsubseteq c')$

The least upper bound of  $lub_{IS_2}(c)$  is the set of all least upper bounds of  $c$  in  $C_2$ .

The rationale of using these approximations is that we can decide whether an entity  $x$  is a member of a class in the private ontology based on its membership in classes of the shared ontology. This decision in turn provides us with an approximate result on deciding whether  $x$  is the result of a query stated in terms of a private ontology, based on the following observation:

- If  $x$  is member of a lower bound of  $c_1$  then it is also in  $c_1$
- If  $x$  is not member of all upper bounds of  $c_1$  then it is not in  $c_1$

In [Selman and Kautz, 1996] Selman and Kautz propose to use this observation about upper and lower boundaries for theory approximation. We adapt the proposal for defining an approximate classifier  $M'$  that assigns members of shared concepts to private ones in the following way:

**Definition 11 (Concept Approximation)** Let  $C_1$  be a set of private concepts,  $C_2$  a set of shared concepts of an agent and  $x$  the member of a shared concepts then for every  $c_1 \in C_1$  we define  $M'$  such that:

- $M'(x, c_1) = 1$  if  $x : \left( \bigvee_{c \in \text{glb}_{IS_2}(c_1)} c \right)$
- $M'(x, c_1) = 0$  if  $x : \neg \left( \bigwedge_{c \in \text{lub}_{IS_2}(c_1)} c \right)$
- $M'(x, c_1) = ?$ , otherwise

Where the semantics of disjunction and conjunction is defined in the obvious way using set union and intersection .

Based on the observation about the upper and lower bounds, we can make the following assertion about the correctness of the proposed approximate classification:

**Theorem 1 (Correctness of Approximation)** The approximation from definition 11 is correct in the sense that:

1. If  $M'(x, c_1) = 1$  then  $x^{\mathfrak{S}} \in c_1^{\mathfrak{S}}$
2. If  $M'(x, c_1) = 0$  then  $x^{\mathfrak{S}} \notin c_1^{\mathfrak{S}}$

Using the definition of upper and lower bounds the correctness of the classification can be proven in a straightforward way [Stuckenschmidt, 2002b].

## 3.2 Queries as Concepts

The results of the last section provide us with the possibility to compute a set of objects that are definitely members of a concept expression and a set of objects that are possibly members of a concept. This approach can directly be used to answer trivial queries that only ask for members of a particular class. We have shown that a slight variation of the mechanism can also be used to approximate Boolean queries over concept names [Stuckenschmidt, 2002a]. In order to compute (approximate) answers for ontology-based conjunctive queries, however, we also have to deal with binary relations in the query expression. In order to cope with relations as well, we use a method for translating conjunctive queries into concept expressions that has been proposed by Horrocks and Tessaris [Horrocks and Tessaris, 2000]. The idea of the approach of Horrocks and Tessaris now to translate the query into an equivalent concept expression, classify this new concept and use standard inference methods to check whether an object is an instance of the query expression. This approach makes use of the fact that binary relations in a conjunctive query can be translated into an existential restriction in such a way that logical consequence is preserved after a minor modification of the A-Box. Details are given in the following theorem.

**Theorem 2 (Role Roll-Up (Horrocks and Tessaris 2000))** *Let  $\langle \mathcal{C}, \mathcal{R}, \mathcal{A} \rangle$  be a description logic knowledge base with concept definitions  $\mathcal{C}$ , relation definitions  $\mathcal{R}$  and assertions  $\mathcal{A}$ . Let further  $R$  be a role,  $C_1$  Concept names in  $T$  and  $a, b$  be individual names in  $A$ . Given a new concept name  $P_b$  not appearing in  $T$ , then*

$$\langle \mathcal{C}, \mathcal{R}, \mathcal{A} \rangle \models (a, b) : R \wedge b : C_1 \wedge \dots \wedge b : C_k$$

*if and only if*

$$\langle \mathcal{C}, \mathcal{R}, \mathcal{A} \cup \{b : P_b\} \rangle \models a : \exists R(P_b \sqcap C_1 \sqcap \dots \sqcap C_k)$$

The transformation of a complete query is more difficult due to the dependencies between the variables that occur in the query expression. In order to keep track of these dependencies during the transformation Horrocks and Tessaris introduce the notion of a query graph.

**Definition 12 (Query Graph (Horrocks and Tessaris 2000))** *The graph induced by a query is a directed graph with a node for every variable and individual name in the query and an directed edge from node  $x$  to node  $y$  for every role term  $(x, y) : R$  in the query.*

The correct transformation of a query into a concept expression depends on the kinds of dependencies between the variables in the query which is reflected in the structure of the query graph. While the approach of Horrocks and Tessaris is more general, we restrict ourselves to queries where the query graph is a (directed) tree and its root node corresponds to the variable we are interested in. In especially, this requires that none of the roles used in the query is declared to be functional and that each constant only appears once in a query. While using this simplification, we would like to emphasize that the translation can be done for unions of conjunctive queries with an arbitrary number of result variables and a very expressive logical language for defining class expressions. Our simplifying assumptions lead to a simple method for transforming a query graph into a concept expression.

**Definition 13 (Query Roll-up (Horrocks and Tessaris 2000))** *the roll-up of a query  $Q$  with query tree  $G$  is a concept expression derived from  $Q$  by successively applying the following rule:*

- *If  $G$  contains a leaf node  $y$  then the role term  $(x, y) : R$  is rolled up according to definition 2. The edge  $(x, y)$  is removed from  $G$*

The result of applying this translation technique to our example query in equation 2.5 is the following expression:

$$\begin{aligned} & (\text{Hotel} \sqcap (\exists \text{liegt} - \text{in} - \text{Ort}.(\exists \text{liegt} - \text{in} - \text{Land}.\{\text{mecklenburg}\})) \sqcap \\ & (\exists \text{hat} - \text{Zimmer}.\leq 25)) \sqcap \\ & (\exists \text{ist} - \text{in} - \text{Schloss}.\{ja\}) \end{aligned} \quad (3.1)$$

As this expression defines a new concept in the overall ontology we can now apply the approximation techniques described in the last section in order to compute the sets of possible and the set of definite answers to the query.

### 3.3 Quality of Approximation

Unfortunately, proving the correctness of the approximation says nothing about the quality of the approximation. In the worst case, the upper and lower boundaries of concepts in the other hierarchy are always  $\top$  and  $\perp$  respectively. In this case the translated query always returns the empty set as result. We were not able to investigate the quality of approximations on theoretical level, however, we can provide some rules of thumb that can be used to predict the quality of an approximation:

**Depths of hierarchies:** The first rule of thumb, we can state is that deeper class hierarchies lead to better approximations. For hierarchies of depth one it is easy to see that we will not be able to find good upper and lower bounds. We can also assume that deeper hierarchies provide finer grained distinctions between concepts that in turn often produce closer approximations.

**Degree of overlap:** Our approach assumes a shared ontology, however, we cannot guarantee that different systems indeed use the same parts of this shared vocabulary. Therefore, the actual overlap of terms used in the existing definitions that are compared is important for predicting the quality of approximations. In general, we can assume that a high degree of overlap leads to better approximations.

Both criteria used in the rules of thumb above strongly depend on the application and on the creator of the corresponding models. At least for the degree of overlap, we can assume that hierarchies that are concerned with the same domain of interest will share a significant part of the vocabulary, thus enabling us to compute reasonable approximations.

In the course of a case study it turned out that in most cases the approximation of concept expressions returns good results, because people tend to share a reasonable number of concept names across different ontologies that provide a basis for creating mappings. These mappings can often be found using stemming and simple string matching. On the other hand, it turned out that it is much harder to come up with reasonable mapping between the relations used in different ontologies leading to a situation where we only have very sparse mappings between these relations. This in turn has a major impact on the quality of approximation applied to conjunctive queries. In fact the lack of mappings between relations often leads to a situation, where answers could not be found, because names of relations in the query were not known in the ontology in the

answering peer. In the next section, we discuss an approach to overcome the problem of sparse mappings between relation names.

## Chapter 4

# Query Relaxation

In the presence of sparse mappings, we face a situation where the descriptions of different peers referring to the same real-world object can be significantly different. In most cases, the descriptions are different in the sense that different relations are used to related same object to other objects in the domain. These relations may refer to the same properties of the object that cannot be matched due to a missing mapping or the set of properties itself used might be different. As a consequence, real-world objects that are meant to be an answer to a query are not returned because their description does not match the query that is formulated using terms from a different ontology. We address this problem by relaxing the query, i.e. by weakening those constraints from the query expression that are responsible for the failure. In order to be useful, this weakening process has to fulfill certain formal properties. In especially, we want to make sure that we do not lose any answers when modifying the query. We can guarantee this using the notion of query subsumption as described by Halevy:

**Definition 14 (Query Containment and Equivalence (Halevy 2001))**

Let  $\mathcal{T} = \langle \mathcal{C}, \mathcal{R}, \mathcal{O} \rangle$  and  $Q_1, Q_2$  conjunctive queries over  $\mathcal{T}$ .  $Q_1$  is said to be contained in another query  $Q_2$  denoted by  $Q_1 \sqsubseteq Q_2$  if for all possible sets of object definitions of a terminological knowledge base the answers for  $Q_1$  is a subset of the answers for  $Q_2$ :  $(\forall \mathcal{O} : \text{res}(Q_1) \subseteq \text{res}(Q_2))$ . The two queries are said to be equivalent, denoted as  $Q_1 \equiv Q_2$  iff  $Q_1 \sqsubseteq Q_2$  and  $Q_2 \sqsubseteq Q_1$

Based on these notions we compute a sequence of queries  $Q_0, \dots, Q_n$  such that the following properties hold:

1.  $Q_0 \equiv Q$
2.  $i < j \implies Q_i \sqsubseteq Q_j$

The intuition behind this approach is to start with the original query and generate queries where each is more general than the one before, i.e. each query following in the sequence returns all results of the previous one, but might



return more results. Our hope is that these new results contain the description of some real-world objects that should be answers, but were not found due to their description.

There are many different ways of making a query more general in order to increase the chance of matching a potential answer. In the following we discuss relaxation heuristics we consider useful for the purpose of query processing in a peer-to-peer setting.

## 4.1 Variable Elimination

The first heuristic is based on the fact that each variable in a conjunctive might fail to match a specific object if the object does not satisfy the constraints. Therefore, a way of increasing the chance of matching the target object in the head of the query is to successively eliminate non-answer variables from the query. In the example query in equation 2.5 for example, we have the variables  $V, W, X, Y$  and  $Z$  where  $X$  is the answer variable. Therefore we can weaken query by eliminating the variables  $V, W, Y$  and  $Z$ . This can be done by removing all conjuncts containing a specific variable from the query expression. It is easy to see that successively removing conjuncts from the query leads to a sequence of queries with the desired properties.

The main question that arises when adopting the variable elimination approach is the order in which the variables should be removed from the query. This order is partially constrained by the dependencies between the different variables. Removing the wrong variable first can break these dependencies and make remaining conjuncts useless. Looking at the example query this would happen if we first removed the variable  $Y$ . In this case the conjunct  $V = mecklenburg$  would be isolated, because the variable  $V$  only occurred in the removed conjuncts that connected it to the answer variable. In order to avoid breaking dependencies when removing conjuncts, we can use the query graph of the query to be relaxed (compare definition 12) as it explicates existing dependencies. In the query-graph dependencies between variables are represented by arcs between nodes. Therefore we have to ensure that the query graphs remains connected when removing the node that represents the variable we want to eliminate. Obviously, this is only the case if we eliminate variables that correspond to leaf nodes in the graph. Figure 4.1 illustrates the successive elimination of the variables  $V, Y, Z$  and  $W$  from the example query, showing the corresponding sequence of query graphs.

## 4.2 Guided Elimination

The major drawback of the variable elimination heuristic as explained so far is the high number of arbitrary choices that still exist in the order of elimination.

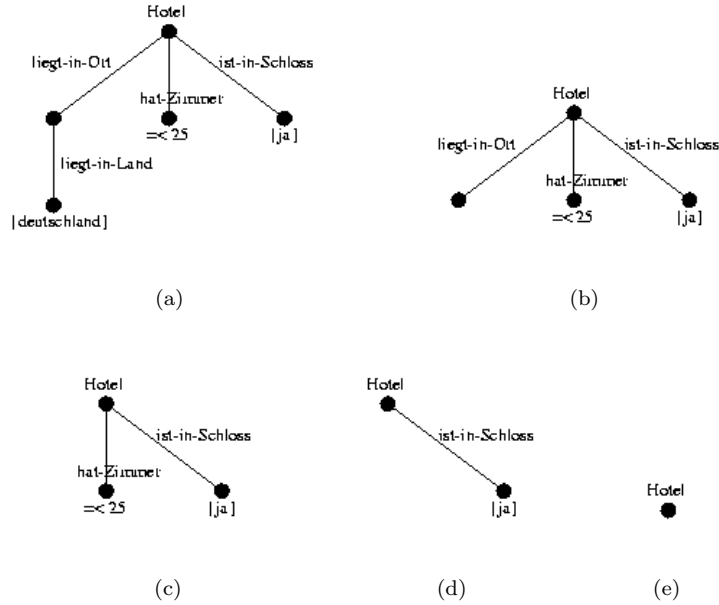


Figure 4.1: A possible sequence of query-graphs

More specifically, whenever the query tree has more than one leaf node, we have no strategy yet to decide which one to eliminate. In general, there are many possibilities for defining ordering heuristics, based on:

1. The nature of the domain
2. The preferences of the user
3. The task to be solved

As our approach does not aim at a specific domain, user or task, we will have to rely on rather general heuristics being aware that they will never be optimal. In our case, the only information we can use to decide on an elimination order is the existence of local mappings that relate the query vocabulary to the shared one that is actually used to compute the answer. The general idea is that we would rather drop conjuncts that represent concepts or relations without a suitable mapping into the shared ontology, because it can never be satisfied by any object classified according to that ontology. We have seen that for the case of concepts, we can often find a suitable approximation even if there is no direct counterpart in the shared ontology. Therefore, we focus on conjuncts representing relations and eliminate such variables first that are constrained by a relational conjunct that has no direct mapping to the shared ontology. The

effect of this strategy is illustrated in the next section where we describe some experiments with approximating concepts and relaxing queries in a case study.

## Chapter 5

# Examples from a Case Study

We performed a first case study in order to validate the methods described in this paper. The case study is based on three different ontologies in the domain of tourism. The ontologies are available in the DAML ontology library ([www.daml.org](http://www.daml.org)) and have been created by independent groups of students at the university of Karlsruhe. All ontologies aim at describing the conceptualization of an internet site that is advertising tourism in north-east Germany. All ontologies contain information about accommodation, tourist attractions and transportation facilities. While sharing these general topics, the different ontologies describe them in a very different way focusing on different parts of the overall domain. We chose these ontologies, because they very closely resemble the situation we expect in a peer-to-peer network, where peers model information about the same domain in different ways.

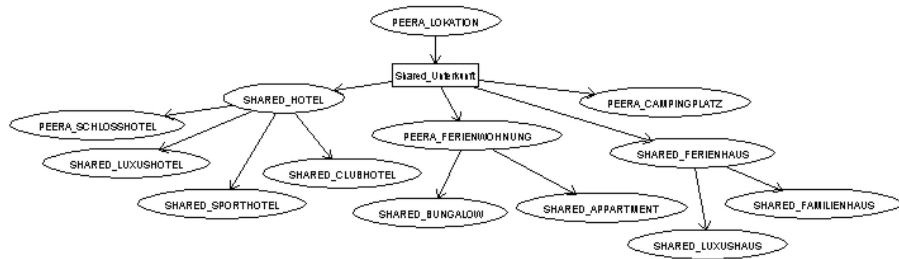
In the course of our case study, we imported the ontologies, each containing about 300 concepts and 50 to 70 relations into an ontology editor using some syntactic transformations. We then analyzed the ontologies and created about 150 obvious mappings, mainly between concepts that have exactly the same name and between concepts where one name is the plural form of the other. Based on these mappings we computed two overlapping concept hierarchies consisting of about 600 concepts each. These hierarchies served as the basis for evaluating our concept approximation and query relaxation techniques. In the following, we describe examples of concept approximation and of query relaxation with respect to this hierarchy.

### 5.1 Concept Approximations

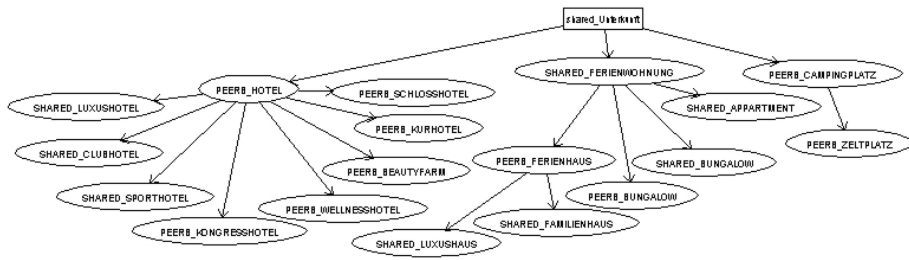
As an example of concept approximation we use the concept 'Ferien-Wohnung' (a flat used as accommodation during holidays). The relevant part of the

overall hierarchy can be seen in figure 5.1. We can see that concepts from private and shared ontologies occur in this part of the hierarchy.

The approximations we are interested in are the direct sub- and super-classes of our example concept that are not from the same ontology. We can see in the figure that these are: 'Bungalow' and 'Apartment'. If we look at the view of PeerB on the world we see that also the concept 'Ferienhaus' (house used during holiday) would fall under this category. While this result is not completely true, because houses are not flats, it still serves the purpose very well, because all of concepts describe accommodations that are reasonable replacements in the case that no flat is available.



(a) Peer A



(b) Peer B

Figure 5.1: The Views of two different peers on the same domain

If we determine the upper approximation of the example concept, we get the general concept 'Unterkunft' (accommodation). Our method now determines all instances of this general concept to be potential members of the example concept. Besides the members of the already mentioned concepts, this also

includes objects that are members of the concepts 'Hotel' and 'Campingplatz' (Camp-Site) in the view of the answering peer B. We see, that these results are still closely related to the example concept, because they are all accommodations mainly used during holiday, however, hotels and Camp-Sites are not really the kind of answer the user would assume to get when asking for a flat. Still, returning hotels and camp-sites as answers to a query for a flat is still better than not returning any result, because the user might want to change her choice in favor of other preferences (e.g. the location).

## 5.2 Query Relaxation

As an example for query relaxation, we take the example query introduced in equation 2.5. If we transform this query in to a concept expression (equation 3.1) and classify it into the overall concept hierarchy of the case study, it end up as a sub-concept of 'Schlosshotel' (castle acting as a hotel). Computing the answer to the query we get an empty set, because there are no instances of 'Schlosshotel' satisfying all properties of the query concept. Using the upper bound, however, we already get the members of the the concept 'Schlosshotel'. If we do not want to rely on this result, we have to analyze the reason for the failure of returning definite answers. Looking at the ontologies in the case study, we see that none of the ontologies except for one the query is based on contains information about the number of rooms of a hotel which makes it impossible to prove that a specific Hotel is an answer to the query. As a response to this observation, we relax the query by removing the restriction on the number of rooms. This leads to a situation, where we already get some definite results, namely those members of the concept 'Schlosshotel' that satisfy the requirement of being in the federal state of Mecklenburg. Note that this provides us with a better result than the use of the upper bound, because we already have a pre-selection of results according to the geographic criterion.

The ability to retrieve relevant information using this second query relied on the fact that the ontology describing the information defines the concept 'Schlosshotel' as the set of all hotels for which the property 'liegt-in-schloss' (is located in a castle) is true. We were able to use this implicit information about the specific relation in order to retrieve information without having an explicit assertion stating that a hotel has this specific property. In a case where the ontology does not contain the necessary information, we would still get no results for the relaxed query, because the property 'liegt-in-Schloss' is not satisfied by any information item. In this case we can again use the upper bound for answering the query, which would now be the concept 'Hotel'. Consequently, we would get all Hotels as potential answers. Again, this result is too general, as we want to preserve at least the geographic constraint. A solution is to further relax the query by removing the 'liegt-in-Schloss' property from the query. The resulting query will match all Hotels in the federal estate of Mecklenburg. Admittedly, this result is not a very good one, however, it resembles the func-

tionality of many current web-based information systems, where lists of hotels can be accessed by selecting a specific area. We would like to stress that our approach leads to more precise answers in most cases, especially if the queries are not too complicated and we only have to fall back to very imprecise results if all other attempts fail.

## Chapter 6

# Conclusions

The existing data integration technology is based on the assumption that it is possible to define a virtual global schema. A query is therefore posed to this schema and then suitably translated in local queries to the information bases being integrated. As hinted in this paper and more extensively argued elsewhere, this approach hardly scales to a highly dynamic P2P network.

In this paper we have proposed a novel approach based on the following key ideas:

1. We shouldn't think in terms of a global schema but, rather, in terms of independent autonomous nodes which, at run time, depending on the query, provide local answers which must then be integrated.
2. In most cases, the system will not be able to provide the best possible answer. It is more an issue of providing an answer which is good enough.

The proposed approach exploits ontologies as a conceptual "high level" schema which allows to hide local implementation details, and to exchange information using the vocabulary defined by the ontologies themselves. In this setting, good enough answers are obtained by posing queries to local ontologies, by allowing queries to be propagated by using inter-ontology mappings, and by using approximation techniques in order to avoid the problem of very sparse mappings.

We are only at the beginning and a lot of work is still to be done before integrating Semantic Web and P2P network technology. We list below some of the open problems we can foresee:

**Foundations.** Once we assume that we have no global schema we must also assume, among other things, that we no longer have a single domain of interpretation (a single set of models). To take this into account we must define a new semantics which allows for multiple interpretation domains, and for mappings which tell how ontology elements, and also domain elements, are mapped. In this framework, queries, query answers, and the mappings defined in this



paper can be formally characterized and various (partial) completeness and correctness results can be provided (characterizing, for instance, to which extent, a query has been answered). Some preliminary ideas in this direction can be found in [YYY].

**Matching.** By matching we mean here the problem of taking two semantically heterogeneous ontologies and finding a mapping (one-to-one, many-to-one, many-to-many) which puts in correspondence terms with similar meaning (according to some "success" metric). This research area, though not mature, has been studied extensively in the database community (see for instance [9]). In the P2P domain further dimensions of complexity exist; for instance: in most cases one would like to approximate real time query answering (while matching is an inherently complex problem), matching should be (possibly) completely automated, due to the strong dynamics of a P2P network, it is hardly to foresee the nodes which will need to be "matched".

**Semantic Routing.** In a P2P network, in most if not all cases, a node has very little knowledge of its peers, and this knowledge becomes obsolete very quickly. Before worrying about how to answer a query (the problem we have mainly dealt with in this paper) a node has to worry about which nodes should be queried. This requires some mechanisms for peer discovery that, in general, will have again to deal with the problem of semantic heterogeneity (each node will advertise itself using some local vocabulary). A preliminary description of a possible solution has been provided in [XXX], which proposes "Interest Groups" as a way of collecting peers having knowledge about same or similar topics and a "locally centralized" mechanism for handling such groups.

**Good enough answers.** We have already given up the hope for complete or, more simply, correct answers. But how do we judge when an answer is good answer? Whether an answer is good enough depends on many things: what the user wants, the status of the network, its connectivity, the topic of the query, and so on. We need to provide techniques for formalising this fact and for measuring what is good enough. This is necessary step for deciding what we can do with the results we get from a P2P network.

# Bibliography

- [Bergamaschi et al., 1999] Bergamaschi, S., Castano, S., and Vincini, M. (1999). Semantic integration of semi-structured and structured data sources. *SIGMOD Records*, 28(1):54–59.
- [Calvanesea et al., 2001] Calvanesea, D., De Giacomo, G., and Lenzerini, M. (2001). Description logics for information integration. In *Computational Logic: From Logic Programming into the Future*, Lecture Notes in Computer Science. Springer Verlag.
- [Chang and Garcia-Molina, 2001] Chang, K.-C. and Garcia-Molina, H. (2001). Approximate query mapping: Accounting for translation closeness. *The VLDB Journal*, 10:155–181.
- [Donini et al., 1996] Donini, F. M., Lenzerini, M., Nardi, D., , and Schaerf, A. (1996). Reasoning in description logics. In Brewka, G., editor, *Principles of Knowledge Representation*, Studies in Logic, Language and Information, pages 193–238. CSLI Publications.
- [Euzenat, 2001] Euzenat, J. (2001). An infrastructure for formally ensuring interoperability in a heterogeneous semantic web. In Cruz, I. F., Decker, S., Euzenat, J., and McGuinness, D., editors, *Proceedings of the First Semantic Web Working Symposium*, pages 345–360, Stanford, CA, USA.
- [Garcia-Molina et al., 1997] Garcia-Molina, H., Papakonstantinou, Y., Quass, D., Rajaraman, A., Sagiv, Y., and Ullman, J. (1997). The TSIMMIS project: Integration of heterogeneous information sources. *Journal of Intelligent Information Systems*, 8(2):117–132.
- [Goasdoué and Rousset, 2000] Goasdoué, F. and Rousset, M.-C. (2000). Rewriting conjunctive queries using views in description logics with existential restrictions. In *Proceedings of International Workshop on Description Logics - DL'2000*, volume 33, pages 113 – 122. CEUR-WS.
- [Gomez-Perez and Corcho, 2002] Gomez-Perez, A. and Corcho, O. (2002). Ontology languages for the semantic web. *IEEE Intelligent Systems*, January/February:54–60.

- [Halevy, 2001] Halevy, A. Y. (2001). Answering queries using views: A survey. *The VLDB Journal*, 10(4):270–294.
- [Horrocks, 1998] Horrocks, I. (1998). The FaCT system. In de Swart, H., editor, *Automated Reasoning with Analytic Tableaux and Related Methods: International Conference Tableaux’98*, number 1397 in Lecture Notes in Artificial Intelligence, pages 307–312. Springer-Verlag, Berlin.
- [Horrocks and Tessaris, 2000] Horrocks, I. and Tessaris, S. (2000). A conjunctive query language for description logic aboxes. In *AAAI/IAAI*, pages 399–404.
- [Lenzerini, 2002] Lenzerini, M. (2002). Data integration: A theoretical perspective. In *PODS*, pages 233–246.
- [Levy, 1999] Levy, A. (1999). Combining artificial intelligence and databases for data integration. In Wooldridge, M. and Veloso, M., editors, *Artificial Intelligence Today: Recent Trends and Developments*, volume 1600 of *Lecture Notes in Computer Science*, pages 249–268. Springer, Berlin.
- [Levy et al., 1996] Levy, A., Rajaraman, A., and Ordille, J. J. (1996). Querying heterogeneous information sources using source descriptions. In *Proceedings of the 22nd International Conference on Very Large Databases, VLDB-96*, pages 251–262, Bombay, India.
- [Madhavan et al., 2002] Madhavan, J., Bernstein, P., Domingos, P., and Halevy, A. (2002). Representing and reasoning about mappings between domain models. In *Eighteenth National Conference on Artificial Intelligence (AAAI’2002)*, Edmonton, Canada.
- [Papakonstantinou et al., 1996] Papakonstantinou, Y., Gupta, A., and Haas, L. (1996). Capabilities-based query rewriting in mediator systems. In *Proceedings of 4th International Conference on Parallel and Distributed Information Systems*, Miami Beach, Flor.
- [Parameswaran et al., 2001] Parameswaran, M., Sursala, A., and Winston, A. (2001). P2p networking: An information sharing alternative. *IEEE Computing*, 34(7).
- [Rahm and Bernstein, 2001] Rahm, E. and Bernstein, P. A. (2001). A survey of approaches to automatic schema matching. *The VLDB Journal*, 10(4):334–350.
- [Selman and Kautz, 1996] Selman, B. and Kautz, H. (1996). Knowledge compilation and theory approximation. *Journal of the ACM*, 43(2):193–224.
- [Stuckenschmidt, 2002a] Stuckenschmidt, H. (2002a). Approximate information filtering with multiple classification hierarchies. *International Journal of Computational Intelligence and Applications*. to appear.

- [Stuckenschmidt, 2002b] Stuckenschmidt, H. (2002b). *Ontology-Based Information Sharing in Weakly-Structured Environments*. PhD thesis, Faculty of Sciences, Vrije Universiteit Amsterdam.
- [Wache et al., 2001] Wache, H., Voegelé, T., Visser, U., Stuckenschmidt, H., Schuster, G., Neumann, H., and S.Huebner (2001). Ontology-based integration of information - a survey of existing approaches. In *Ontologies and Information Sharing*, number 47, pages 108–117, Seattle, USA.