Finding and Integration of Information - A Practical Solution for the Semantic Web -

Ubbo Visser and **Gerhard Schuster**[1]

Abstract. If we believe the numerous publications concerning intelligent approaches for better information retrieval from the WWW the Semantic Web is already alive. However, the nature of most of the approaches is more theoretical. One major outcome of the research being undertaken over the last few years in the area of artificial intelligence for the Semantic Web is the benefit of using ontologies for content-based information retrieval. This led to a number of systems that provide user interfaces and intelligent reasoning services to access and integrate information sources (e.g. Ontobroker, SHOE, OntoSeek, BUSTER). This paper deals with a practical solution for finding and integrating information from the Web. Since some of the ideas of our BUSTER system are already known we focus on two issues: we introduce the Comprehensive Source Description (CSD), a necessary description for information sources that allows extra services such as integration or translation and a new feature that allows a combined search for concepts at a certain location, introducing the *concept@location* query. We discuss implementation issues and provide an example for better understanding.

1 INTRODUCTION

The Internet as de facto biggest information source electronically available consists of a vast amount of data, which are mainly loosely structured. Mostly, these data belong to proprietary systems, which are not build for interoperability in the first place. With the comprehensive networking it is nowadays possible to link the items in the network together. Thus, there is a need for tools that are able to find, access, and integrate the information sources. The main obstacles are schematic and semantic heterogeneity problems, which are thoroughly discussed in various papers [6, 14, 25]. Over the last decade several approaches with regard to intelligent information integration have been proposed (e.g. IM, SIMS, OBSERVER, COIN; see [25] for an overview). The majority of these systems provide representation mechanisms for ontology-based content explication. The systems mainly use some kind of description logics (e.g. OIL). The main reason behind this is the option to explicitly describe concepts of an application domain using a language that provides formal semantics. Lately, this general approach of using ontology-based systems for information integration has been widely accepted [8].

Ontologies became a popular research topic in the 90ies and are still the focus of researchers in the artificial intelligence area. There is still a need for more fundamental research in various areas: the role of ontologies, acquisition of ontologies, semantic mapping and translation to name only a few topics. This may only be one reason why ontology-based systems are mainly theoretical approaches with some prototypical front-ends. A new article in a trendy computer magazine [27] states that there are numerous publications with respect to the Semantic Web but there are only a few applications available. There is a need for practical solutions. The BUSTER system is a contribution for this demand as it provides means for ontology-based search and integration.

In this paper we discuss the BUSTER approach focussing on the description of information sources and describe our prototypical implementation. We introduce a new feature of the system, namely the option to search for concepts using a terminological reasoning service and to search for locations using a spatial reasoning service. A combination of both leads to a new query type *concept@location*, e.g. "Are there land cover sources available that cover Lower-Saxony?" or searching for "suppliers for product X in region Y".

2 APPROACH

The Bremen University Semantic Translator for Enhanced Retrieval (BUSTER), a middleware based also on ontologies, has been developed at the Center for Computing Technologies. BUSTER is based on the hybrid ontology approach, i.e. it can access more than one ontology and integrate them. The only restriction is that there is a common vocabulary the ontologies are based on. Schuster and Stuckenschmidt [21] describe a method that leads to a common vocabulary using known but domain dependent thesauri.

The concept view of the system is shown in figure 1. It shows the query phase on the right hand side and the acquisition phase on the left hand side. Since the description of information sources with meta data is crucial we focus on the *Comprehensive Source Descriptions* (CSD), located at the site of the data source or service, and formalized in XML/RDF format. A thorough description about concept of BUSTER can be found in [22].

Comprehensive Source Description

In order to describe existing data metadata have to be used. Hence, we have to find an eligible language for the description. Over the last decade numerous meta data formats have emerged (e.g. Dublin Core, ISO/TC211). A good overview about existing meta information systems can be found in [23]. Since we are not dependent on any specific domain, in fact we would like to use a general way to describe the data, we use the Dublin Core Element Set, version 1.1 as a de facto basis for our CSD. The definitions utilize a formal standard for the description of metadata elements. The authors claim that the formalization helps to improve consistency with other metadata communi-

[1] Center for Computing Technologies, Universität Bremen, Universitätsallee 21-23, D-28359 Bremen, Germany, email: {visser|schuster}@tzi.de

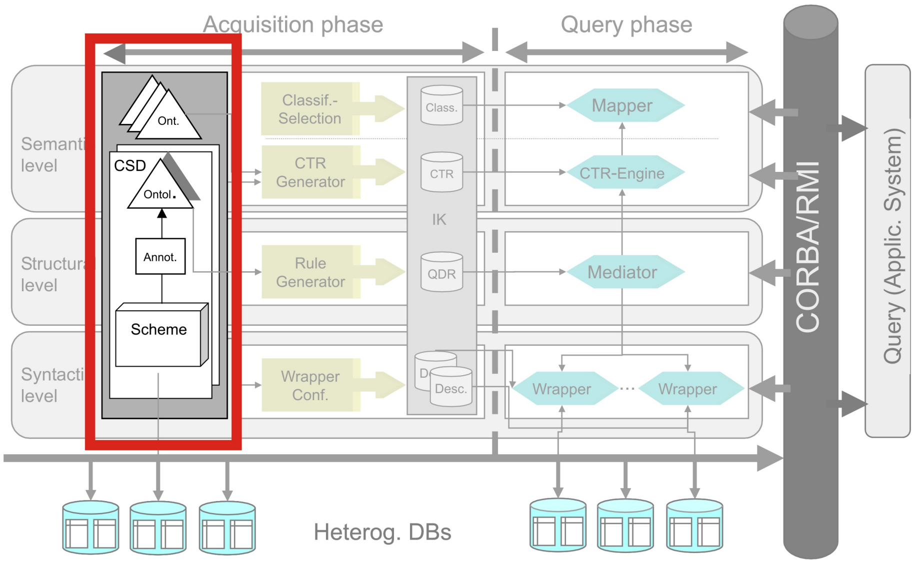

Figure 1. BUSTER: concept view

ties and enhances the clarity, scope, and internal consistency of the Dublin Core metadata element definitions.

However, some of the given elements are not sophisticated enough in their expressivity (e.g. the relation element) or lack formal semantics (e.g. description element). Thus, there is a need for additional qualifiers for those elements, which are described in a language that provides formal semantics (e.g. DAML, OIL, SHIQ). We can use this kind of description logics to encode additional features. We use the RDF(S) syntax if possible to ensure a wide acceptance with respect to accessibility and usability. We then refer to explicit ontologies available on the WWW. The following Dublin Core elements are refined for our CSD:

- Coverage: Since there is no further distinction between spatial and temporal coverage, this element has to be refined.
  - Spatial: The recommended best practice from DCMI is to select a value from a controlled vocabulary and that, where appropriate, named places or time periods be used in preference to numeric identifiers such as sets of coordinates or date ranges. Examples are *DCMI Point* to describe a point in space using its geographic coordinates, *ISO 3166* a code for the representation of names of countries, *DCMI Box* that identifies a region of space using its geographic limits. The last recommendation is *TGN*, the GETTY Thesaurus of Geographic Names (see http://shiva.pub.getty.edu/tgn_browser/). We decided on the latter because the use of place names is more intuitive and therefore more valuable with respect to users on the WWW.
  - Temporal: The recommend best practice here is to use one of the two following encoding schemes: *DCMI Period*, a specification of the limits of a time interval, and *W3C-DTF*, the W3C encoding rules for dates and times - a profile based on ISO 8601 (see also: http://www.w3.org/TR/NOTE-datetime). We use the latter since the main reason to have this CSD is to describe information sources on the WWW.
- Description: Description may include but is not limited to: an abstract, table of contents, reference to a graphical representation of content or a free-text account of the content. The semantics of this kind of representation are limited with regards to machine readable meaning of the content. Hence, we restrict the description to a formal description logic, namely DAML+OIL or SHIQ. The vocabulary used to describe this A-Boxes has to be one of the vocabularies used in the ”relation” element.
- Relation: The qualifiers that refine the relation element as recommended by DCMI is limited. Therefore, we need to extend these qualifiers by references that also point to ontologies, gazetteers or thesauri. A relation is described as a XML namespace describing the URI of the corresponding vocabulary and a prefix to mark terms from this vocabulary.
- Subject: The qualifiers recommended by DCMI for the subject element contain common lists of keyword from various sources (e.g. the Library of Congress Subject Headings, Medical Subject Headings, Universal Decimal Classification). In BUSTER, we use the subject element accordingly, it remains a list of significant keywords to describe the information source but the keywords have to be chosen from a controlled vocabulary referred by the relation element.
- Rights: Despite the intellectual property rights we also have to consider access rights for special user groups. In the moment, there is no further specification.

Figure 2 shows an extract from a typical CSD, a CSD for a data set concerning land use in Lower-Saxony, Germany in this particular case. We only show the relevant parts according to the refined elements mentioned above. The subject contains links to a ”topic-area” described in the general CSD ontology and some concepts concerning the content of the topic-area described in the GEMET ontology. The ”description” element consists of two additional properties of that information source (a) the fact the data set consist of a Bessel-ellipsoid from 1841, which is described in a geodesic ontology, and (b) the meaning of the attributes of the underlying relational table. One might think that this is additional modeling effort for no good

```xml
 <?xml version="1.0" encoding="ISO-8859-1" ?>
- <rdf:RDF
  …
  <dc:title>Database for land use in southern Lower Saxony, Germany</dc:title>
  <dc:type>dataset</dc:type>
  <dc:source>http://www.tzi.de/buster/data/csd/clc.dbf</dc:source>
- <dc:subject>
    <csd:topic-area rdf:resource="gemet:land-use" />
    <csd:topic-area rdf:resource="gemet:land-use-classification" />
    <csd:topic-area rdf:resource="gemet:landscape-utilisation" />
  </dc:subject>
- <dc:description>
    <geodesy:reference   rdf:resource="geodesy:Bessel-ellipsoid-1841" />
    <csd:table-name      rdf:resource="corine:clc" />
    <csd:table-attribute rdf:resource="csd:row_number" />
    <csd:table-attribute rdf:resource="csd:identifier" />
    <csd:table-attribute rdf:resource="corine:nomenclature" />
    <csd:table-attribute rdf:resource="geodesy:Bessel-ellipsoid-1841" />
    <csd:table-attribute rdf:resource="csd:tknr" />
  </dc:description>
- <dc:relation>
    <csd:reference alias="csd"     source="csd.rdfs" />
    <csd:reference alias="gemet"   source="gemet.rdfs" />
    <csd:reference alias="corine"  source="corine.rdfs" />
    <csd:reference alias="geo"     source="germany.xml" />
    <csd:reference alias="tgn"     source="tgn.daml" />
    <csd:reference alias="geodesy" source="geodesy-ontology.rdfs" />
  </dc:relation>
- <dc:coverage>
    <geo:state  rdf:resource="tgn:Niedersachsen" />
    <geo:region rdf:resource="geo:Northwest-Germany" />
  </dc:coverage>
- <dc:creator>
+ <dc:rights>
…
  </rdf:RDF>
```

Figure 2. Extract of a typical Comprehensive Source Description (CSD)

reason but we are now able to enable additional services such as automatic translation processes between catalogue systems as described in [17].

Based on the metadata provided by the CSDs and appropriate qualitative terminological (conceptual ontologies) and spatial models (spatial ontologies), BUSTER supports integrated queries of the type *concept@location*. These type of queries are described in the next section.

3 IMPLEMENTATION

The prototype of the BUSTER is based on an open server-client architecture, and can be divided into two main parts: the so-called BUSTER-cluster on the server side and a BUSTER client.

3.1 Architecture

BUSTER clients can be started as local applications or as java applets in a standard browser supporting Java Swing. The BUSTER client provides an ontology-driven user interface to specify queries and to present the results of the retrieval. Additional services such as automatic translation process will be made available dependent on the result and if applicable. The communication between the clients and the cluster is implemented via Remote Method Invocation (RMI).

The BUSTER cluster comprises several modules relevant for intelligent querying and semantic translation purposes: a BUSTER server, a database for CSDs and available domains, a web server, and spatial and logical reasoning modules (see figure 3). Examples for the latter available on the WWW are the FaCT system provided by the University of Manchester [13] and the RACER system provided by the University of Hamburg [10]. These modules are within the BUSTER cluster to fit the minimum requirements for terminology and spatial queries, but its open architecture allows to use arbitrary services for reasoning, translation or other tasks if needed.

An Apache web server provides the platform for the applets. The server handles client queries depending on the users selection. It controls the process of the query *(concept@location)* by retrieving domain specific information from a SQL-database via JDBC interface, downloading distributed CSDs and knowledge bases, and triggering reasoning services within or outside the BUSTER cluster.

3.2 Queries

Once an information source has been annotated with all the information needed, complex queries can be made to the BUSTER system. BUSTER is based on terminological ontologies that have been modelled in advance. The system allows two different types of queries, a terminological and a spatial query. We introduced the query *concept@location* however, it is possible to submit a terminological query alone without spending time on the spatial part. Also it is possible to submit a spatial query on its own.

3.2.1 *Terminological Queries*

The terminological query can be divided into two parts, namely a simple concept query and a defined concept query.

Simple Concept Query The user chooses one terminology (ontology) depending on current domain, e.g. ”installation supplies”. These terminologies are registered at the BUSTER server. The user can then select on one of the concepts of the taxonomy that fits his query best (e.g. ”installation pipe”). The BUSTER server receives the query and integrates the known terminologies for the current domain by loading them into the connected reasoner. This is possible, because every terminology is annotated with a common vocabulary (hybrid approach, see section 2). After re-classification, all sub-concepts (children) of the query concepts form the result.

Defined Concept Query According to the domain, the user chooses a query-template provided by the BUSTER server. This template contains attributes (slots) and values (filler) from the common vocabulary. The user-interface is ontology-driven, which simply means that the available attributes and fillers are automatically loaded and presented dynamically. This way the user can’t make a mistake, e.g. using unknown terms. The user defines his query by selecting

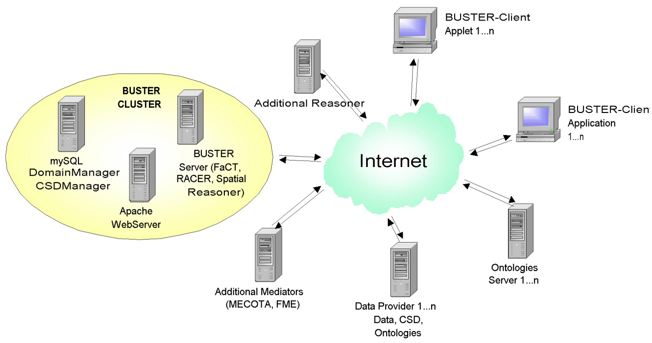

Figure 3. BUSTER: system architecture

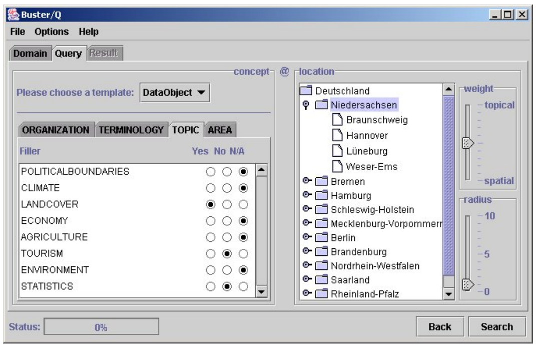

Figure 4. Example for a defined concept and spatial query

reasonable values for the given attributes. 'Yes' specifies the occurrence of the related filler, 'no' prohibits the occurrence and 'n/a' is chosen, if the value does not matter. Figure 4 shows an example of a defined concept query. The user is interested in information about the land cover of Lower-Saxony, a state of Germany. He chooses an appropriate query-template "DataObject", which provides attributes and values for the definition of information sources like databases. He selects the value "landcover" for the attribute "topic". He is not interested in information sources that deal with tourism or statistics. No statements are made about the other values. The filled query-template is translated into a logical term. During the query process all CSDs related to the current domain are parsed for the subject-tag. Each subject references to a namespace, which points to an ontology that contains a concept description of the subject term. These ontologies are then downloaded from ontology servers available on the WWW, are merged with the defined concept query and transferred into available inference machines. After re-classification, all sub-concepts (children) of the query concepts form the result.

In case of a simple concept query, the user has to choose a specific terminology. This makes the query simpler to understand for a user, but it assumes that the user knows at least one terminology or concept from the hierarchy. Simple concept queries are fast, but not always expressive enough. To overcome these problems one could use the defined concept query. On the base of the given common vocabulary the user is able to define a concept that fits his vision of a concrete concept. A defined query is more complex to build, but it is much more unrestricted.

3.2.2 *Spatial Queries*

A user-friendly and, from a cognitive perspective, sound method to specify spatial queries as well as to index data sources and services is the use of placenames. Placenames are typically organized in gazetteers [12, 18]. Schlieder et al. [19, 20] propose an extension to gazetteers in the form of placename structures based on qualitative spatial models. A placename structure can be seen as a hierarchical tree, where the nodes of the tree represent well known name descriptors for geographic features, and the edges reflect their binary part-of relations. These models, or spatial ontologies, use graph representations of hierarchically organized polygonal tessellations as a basis to reason about the spatial relevance of one placename with respect to another. In a qualitative spatial model tree leaves corresponding to nodes of the used connection graph represent the tessellation (see figure 5). Spatial relevance, a combined evaluation of partonomic and neighborhood relations between placenames, is computed by calculating the horizontal and vertical (or hierarchical) graph-theoretical distances.

In BUSTER the user is able to select a specific spatial ontology to initialize a spatial query. In our example the spatial model of Germany is selected. By selecting a placename (e.g. "Niedersachsen"), the user defines the target area of the spatial query. Using the selected spatial ontology, the spatial reasoner integrated in the BUSTER server evaluates the query and computes a list of placenames that are spatially relevant to the target placename. The user is able to parameterize the query by adjusting weight sliders for horizontal and vertical relevance. The example query is configured to find only information sources that are vertical relevant, like sources annotated with "Niedersachsen" or "Hannover" (district of lower-saxony).

3.2.3 *Combined Spatio-Terminological Queries*

BUSTER combines both lists, the list of relevant concepts, and the list of spatially relevant placenames, into one database query. This database query is applied to the BUSTER CSD database. The result is a weighted list of data sources and services matching both the terminological and the spatial query. Figure 6 shows the result of our combined query example. The data source found is an excel sheet with data classified by the Corine land cover nomenclature [4]. Relevant information as well as applicable services from the retrieved CSD are presented. As for the additional service, the user can choose the context translation from Corine to ATKIS [1].

4 Related work

Ontobroker [5] is a well known approach that relies on a single ontology for a group of web users. Therefore, both the data providers and the users have complete access and knowledge to all the concepts described in the ontology. Ontobroker is tailored to homogenous Intranet applications, e.g. for knowledge management within companies. Ontobroker relies on F-Logic and offers therefore advanced inference possibilities. KAON, the KArlsruhe ONtology and Semantic Web infrastructure [3] provides a general three-tier conceptual architecture, which consists of a client layer, a management layer, and a storage layer. The idea of this infrastructure based on RDF and ontologies is to provide services for advanced Semantic Web applications. KAON is a growing family of tools for engineering, discovery, management, and visualisation of ontologies.

OntoSeek [9] is designed for content-based information retrieval from online yellow pages and product catalogues. The retrieval techniques are based on lexical conceptual graphs and large linguistic ontologies (Sensus, WordNet). The basic architecture is similar however, BUSTER uses JAVA applets running in an arbitrary browser on

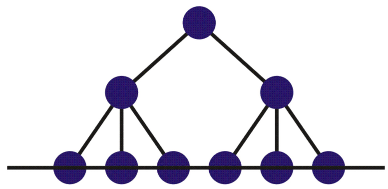

Figure 5. Example for placename structure. The nodes on thee base line represent the tessellation whereas the nodes above represent the placenames

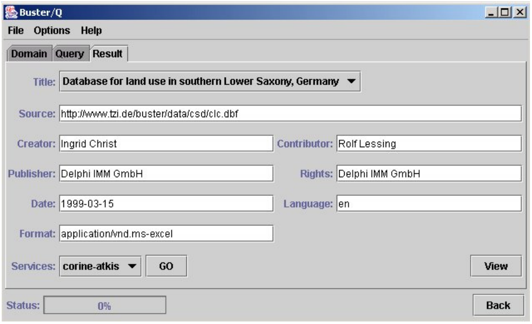

Figure 6. Result for a defined and spatial query

an arbitrary OS. The main difference lies in the expressiveness capabilities of the ontology representation language. In BUSTER we use a more expressive description logic to describe concepts. Another major difference is the possibility to use further services such as a translation service between catalogue systems (if applicable) or a combined search for concepts at locations.

The SHOE Search Tool [11] allows a user to access a SHOE knowledge base by submitting structured queries. This query corresponds to the defined query in our system. The result is presented in a separate window and the user can doubleclick the found URIs to open the corresponding documents. The main differences with respect to BUSTER are the use of ontologies, the query service and other features such as translation services. In Buster, we can use several ontologies for one query, we are able to combine terminological and spatial search, and we can adopt additional mediators for further services shown in the result window.

The Information Manifold (IM) system [15] implements a client with a knowledge base for organizing and querying Internet information sources. The knowledge base contains a rich domain model that enables the description of properties of the information sources. The language used is based on a combination of Horn rules and concepts from the CLASSIC description logic [2]. In contrast to BUSTER IM is based on a single ontology approach using one global ontology. This approach can be applied to integration problems where all information sources to be integrated provide nearly the same view on a domain. If one information source has a different view on a domain, e.g. by providing another level of granularity, finding the minimal ontology commitment becomes a difficult task [7]. Another difference is the restriction of the IM system to only use database sources whereas BUSTER is also able to process other information sources such as XML-based sources.

A major difference between BUSTER and all other mentioned systems is the ability of BUSTER to combine both terminological reasoning and spatial reasoning.

5 CONCLUSION

We proposed a practical solution for finding information sources that have been annotated with metadata and offering additional services for processing the underlying data. We introduced the Comprehensive Source Description (CSD), a necessary description for information sources that allows extra services such as integration or translation. We also proposed a new feature that allows a combined search for concepts at a certain location, introducing the *concept@location* query.

We have seen that *interoperability between terminologies* is possible if we use the hybrid approach [25]. With our proposed practical solution we claim that *system interoperability* is now also feasible. The CSDs are flexible enough to annotate information sources providing additional knowledge to offer extra services. One important extra service could be a translation between different catalogue systems, which we already implemented and introduced elsewhere [17]. A definite drawback is the fact that there is an additional modeling effort. We think that there is a need for automatic annotation facilities. Promising new approaches are Text-To-Onto [16] or the MESA tool [24] for the construction of ontologies. The open architecture of our approach allows the use of additional mediators such as the Feature Manipulation Engine or MECOTA [26].

We think that the *concept@location* query with the included options to regulate both the spatial and the thematic distances are valuable for the Semantic Web. More than 80% of all the data available have an spatial context as we know and we think that our approach is a promising step in the right direction.

REFERENCES

[1] AdV, *Amtliches Topographisch-Kartographisches Informationssystem ATKIS*, Landesvermessungsamt NRW, Bonn, 1998.

[2] Ronald Brachman, Deborah McGuiness, Peter Patel-Schneider, Lori Resnik, and Alexander Borgida, 'Living with classic: When and how to use a kl-one-like language', in *Principles of Semantic Networks: Expo-rations in the Representation of Knowledge*, ed., John Sowa, 401–456, Morgan Kaufman, San Mateo, CA, (1991).

[3] Stefan Decker, Michael Erdmann, Dieter Fensel, and Rudi Studer, 'Ontobroker: Ontology based access to distributed and semi-structured information.', in *Proceedings of Database Semantics 8*, ed., R. Meersman, Semantic Issues in Multimedia Systems., pp. 351–369. Kluwer Academic Publisher, Boston, (1999).

[4] EEA, 'Corine land cover', technical guide, European Environmental Agency, ETC/LC, European Topic Centre on Land Cover, (1997-1999 1997-1999).

[5] Dieter Fensel, Stefan Decker, M. Erdmann, and Rudi Studer, 'Ontobroker: The very high idea', in *11. International Flairs Conference (FLAIRS-98)*, Sanibal Island, USA, (1998). AAAI Press.

[6] Cheng Hian Goh, *Representing and Reasoning about Semantic Conflicts in Heterogeneous Information Sources*, Phd-thesis, MIT, 1997.

[7] Thomas R. Gruber, 'Toward principles for the design of ontologies used for knowledge sharing?', *International Journal of Human Computer Studies*, **43**(5/6), 907–928, (1995).

[8] Michael Grüninger and Mike Uschold, 'Ontologies and semantic integration', in *Software Agents for the Warfighter*, Institute for Human and Machine Cognition (IHMC), University of West Florida, (2002). in preparation.

[9] Nicola Guarino, Claudio Masolo, and Guido Vetere, 'Ontoseek: Content-based acess to the web', *IEEE Intelligent Systems*, **14**(3), 70–80, (1999).

[10] Volker Haarslev and Ralf Möller, 'High performance reasoning with very large knowledge bases', in *International Joint Conferences on Artificial Intelligence (IJCAI)*, ed., Bernhard Nebel, volume 1, pp. 161–166, Seattle, WA, (2001). Morgan Kaufman.

[11] Jeff Heflin and James Hendler, 'A portrait of the semantic web in action', *IEEE Intelligent Systems*, **16**(2), 54–59, (2001).

[12] Linda L. Hill, 'Core elements of digital gazetteers: placenames, categories, and footprints', in *ECDL 2000*, eds., J. Borbinha and T. Baker, Research and Advanced Technology for Digital Libraries, pp. 280–290, Lisbon, Portugal, (2000).

[13] I. Horrocks, 'Fact and ifact', in *Proceedings of the International Workshop on Description Logics (DL'99)*, eds., P. Lambrix, A. Borgida, M. Lenzerini, R. M ö ller, and P. Patel-Schneider, 133–135, CEUR-Workshop Proceedings at http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS, (1999).

[14] Won Kim and Jungyun Seo, 'Classifying schematic and data heterogeinity in multidatabase systems', *IEEE Computer*, **24**(12), 12–18, (1991). problem classification of semantic heterogeneity.

[15] T. Kirk, A. Y. Levy, Y. Sagiv, and D. Srivastava, 'The Information Manifold', in *Information Gathering from Heterogeneous, Distributed Environments*, eds., C. Knoblock and A. Levy, Stanford University, Stanford, California, (1995).

[16] Alexander Maedche and Steffen Staab, 'Ontology learning for the semantic web', *IEEE Intelligent Systems*, **16**(2), 72–79, (2001).

[17] Holger Neumann, Gerhard Schuster, Heiner Stuckenschmidt, Ubbo Visser, and Thomas Vögele, 'Intelligent brokering of environmental information with the buster system', in *International Symposium Informatics for Environmental Protection*, eds., Lorenz M. Hilty and Paul W. Gilgen, volume 30 of *Umwelt-Informatik Aktuell*, pp. 505–512, Zürich, Switzerland, (2001). Metropolis.

[18] W.-F. Riekert, 'Erschließung von fachinformationen im internet mit hilfe von thesauri und gazetteers', in *Management von Umweltinformationen in vernetzten Umgebungen, 2nd workshop HMI*, eds., C. Dade and B. Schulz, Nürnberg, (1999).

[19] Christoph Schlieder, Thomas Vögele, and Ubbo Visser, 'Qualitative spatial representation for information retrieval by gazetteers', in *Conference of Spatial Information Theory COSIT*, volume 2205 of *Spatial Information Theory: Foundations of Geographic Information Science*, pp. 336–351, Morrow Bay, CA, (2001). Springer.

[20] Christoph Schlieder and Thomas Vögele, 'Indexing and browsing digital maps with intelligent thumbnails', in *accepted at Spatial Data Handling 2002 (SDH'02)*, Ottawa, Canada, (2002). Springer. to appear.

[21] Gerhard Schuster and Heiner Stuckenschmidt, 'Building shared terminologies for ontology integration', in *Künstliche Intelligenz (KI)*, pp. 105–123, Wien, (2001).

[22] Heiner Stuckenschmidt, Holger Wache, Thomas Vögele, and Ubbo Visser, 'Enabling technologies for interoperability', in *Workshop: Information Sharing: Methods and Applications at the 14th International Symposium of Computer Science for Environmental Protection*, eds., Ubbo Visser and Hardy Pundt, volume 20, pp. 35–46, Bonn, (2000). TZI.

[23] Ubbo Visser, Heiner Stuckenschmidt, Holger Wache, and Thomas Vögele, 'Using environmental information efficiently: Sharing data and knowledge from heterogeneous sources', in *Environmental Information Systems in Industry and Public Administration*, eds., Claus Rautenstrauch and Susanne Patig, 41–73, IDEA Group, Hershey, USA & London, UK, (2001).

[24] Holger Wache, Thorsten Scholz, Helge Stieghahn, and B. König-Ries, 'An integration method for the specification of rule-oriented mediators', in *International Symposium on Database Applications in Non-Traditional Environments (DANTE'99)*, eds., Yahiko Kambayashi and Hiroki Takakura, pp. 109–112, Kyoto, Japan, (1999).

[25] Holger Wache, Thomas Vögele, Ubbo Visser, Heiner Stuckenschmidt, Gerhard Schuster, Holger Neumann, and Sebastian Hübner, 'Ontology-based integration of information - a survey of existing approaches', in *IJCAI-01 Workshop: Ontologies and Information Sharing*, eds., Asuncion Gómez Pérez, Michael Grüninger, Heiner Stuckenschmidt, and Mike Uschold, pp. 108–117, Seattle, WA, (2001).

[26] Holger Wache, 'Towards rule-based context transformation in mediators', in *International Workshop on Engineering Federated Information Systems (EFIS 99)*, eds., S. Conrad, W. Hasselbring, and G. Saake, Kühlungsborn, Germany, (1999). Infix-Verlag.

[27] Cai Ziegler, 'Deux ex machina', *c't - magazin für computer technik*, **2002**(6), 132–137, (2002).