Context Modeling and Transformation for Semantic Interoperability

Heiner Stuckenschmidt, Holger Wache
Intelligent Systems Group
Center for Computing Technologies
University of Bremen
Email: {heiner, wache}@tzi.de

Abstract

Mediators are middleware components that provide a flexible integration of several heterogeneous information systems. But current approaches ignore that each information has to be considered in its context. Besides the integration the information has to be also converted from its source context into the context of the integrated view in order to achieve interoperability at the semantic level. Two different methods for context transformation are discussed: a rule–based and a classification–based context transformation. We argue that the rule-based context transformation approach does only provide solutions for rather simple transformation tasks. We present an alternative approach that is based on sophisticated context modeling using ontologies based on description logics. We show that this approach supports context transformation based on subsumption reasoning offering reasonable results. Both methods are explained with the help of a real world example coming from the geographical information system domain. We also point out to some open problems that might motivate an extension or modification of existing techniques in the course of further research.

1 Introduction

Mediators [8] are middleware components that provide a flexible integration of several information systems such as database management systems, geographical information systems, or the world wide web. A mediator combines, integrates, and abstracts the information provided by the sources [22] tackling the same problems which are discussed in the federated database research area, i.e. *structural heterogeneity* (schematic heterogeneity) and *semantic heterogeneity* (data heterogeneity) [15]. Structural heterogeneity means that different information systems store their data in different structures. Semantic heterogeneity considers the content and its semantics of an information item. In rule–based mediators [9], rules are mainly designed in order to reconcile structural heterogeneity. Discovering semantic heterogeneity problems and their reconciliation play a subordinate role. But for the reconciliation of the semantic heterogeneity problems, the semantical level also has to be considered [12, 7, 14]. *Contexts* are one possibility to describe the semantical level. A context [13] contains "meta data relating to its meaning, properties (such as its source, quality, and precision), and organization" [17]. A value has to be considered in its context and may be transformed into another context (so–called *context transformation*).

Aim and Scope of the Paper: In the paper we argue that an explication and transformation of context knowledge is required in order to achieve interoperability at the semantic level [5]. Thereby our scope is not on schema integration, another part of the problem. Our approach addresses the problem of preserving the meaning of single attributes in a database: Different databases might use different terminologies or scales for the values of attributes that have been identified as containing the same kind of information by a schema integration step. We briefly review the mediator system MECOTA [21, 20]. We discuss the basic integration mechanism of the system including a rule-based approach to context transformation. We show how the system can be used to integrate the example data-sets. We argue that the rule-based context transformation approach does only provide solutions for rather simple transformation tasks. We present an alternative approach that is based on sophisticated context modeling using ontologies based on description logics [16, 2]. We show that this approach supports context transformation based on subsumption reasoning offering reasonable results for our example data-set (cf. [5]). Thereby our aim is not to develop new reasoning methods and techniques, but to show that well known approaches from the fields of knowledge representation and reasoning can be used to solve this special problem. We also point out to some open problems that might motivate an extension or modification of existing techniques in the course of further research.

An Integration Problem: We illustrate the need for context modeling and transformation by a real–world example which also serves to illustrate our approach. Two sources — `CORINE` and `ATKIS` — provides geological information. For simplicity we assume that both sources store their data in a relational database.

The first source `CORINE` [6] stores its data in two tables[1]. The first table is called `clc_ns2`. Every entry represents one geological item. `clc_ns2` contains the attributes `CLC_NS2_ID` (identifier), `AREA` (surface in ha), and `NS` (classification). Especially the last attribute `NS` refers to catalog, wherein all items are classified. In `CORINE`, the catalog contains more than 64 concepts. The second table `clc_ns2_pol` stores polygons describing the area of an item. The attributes are `CLC_NS_ID` (reference to `clc_ns2`), `VERT_ID` (identifier of a vertices), `RWERT` (x–coordinate of the vertices), `HWERT` (y–coordinate of the vertices), and `NEXT_V_ID` (identifier of the following vertices).

In the second source `ATKIS` [1] a geological item is stored in one table `atkisf` with the attributes `id`, `rechts` (x–coordinate of the first vertice), `hoch` (y–coordinate of the first vertices), `fl` (surface in m^2), and `folie` (classification). Analogously to `CORINE` the last attribute `folie` refers to a classification catalog containing more than 250 terms. But the catalogs of `CORINE` and `ATKIS` are different. Further, both catalogs underly different conceptualizations.

[1]For readability reasons the tables of both sources are simplified. Some attributes are omitted.

The task of this example is that the data of `CORINE` database has to be converted in the `ATKIS` database. Of course, this transformation can be viewed as a special case of an integration task demonstrating all the problems which can be occur. Besides the obvious structural heterogeneity problems, the main problem relies on the reconciliation of the semantic heterogeneity: both geological information sources classify the common areas in different catalogs. During the integration/transformation the classification of the `CORINE` has to be converted into the `ATKIS` catalog. Moreover, the surface has to be converted according to their different currencies. Both conversions are the challenge for the semantic integration and are handled by the both kinds of context transformation.

2 MECOTA - A Rule-Based Mediator

First we focus the reconciliation of structural heterogeneity problems. Integration rules guide the mapping of the information sources to the integrated view provided by the mediator [21]. The rules base on a object–centered description of each information source. Each information item, i.e. an object, a table of a relational database, or simply a number, is encapsulated by *templates*.

A template is an fifth-ary predicate:

$$T = \langle name, context, type, value \rangle @source$$

A template has a *name*, a *context* addressing the semantics of the concept (see below), a *type* determining the data type, the *value* referencing the information item itself, and the last identifier *source* denoting which source the template belongs to. The value can be a simple value, e.g. a number, or a string, or a list of attributes. An *attribute* consists of a name and a template the attribute refers to. In the last case the type is `complex`. In case of simple values the type slot contains the basic data type.

Templates with attributes can represent tables (relations) in a relational data structure model. The attributes of the template are the attributes of the relation. The value of the template attributes are templates encapsulating the basic data types of the relation attributes[2]. An example template for the `ATKIS` table is given below. The template contains variables and therefore describes a set of instances found in the database.

```
<atkisf,?LA,complex, {
  id ->     <id,?LI,string,?ID>,
  rechts -> <rechts,?LR,real,?X>,
  hoch ->   <hoch,?LH,real,?Y>,
  fl ->     <fl,?LS,real,?S>,
  folie ->  <folie,?LC,int,?C>
  }>@ATKIS
```

[2]For readability reasons the source is omitted in the nested templates

An Integration Rule (IR) is designed for the combination and integration of information and primarily for the reconciliation of structural heterogeneity problems. The information is collected according to the templates in the body of the rule. Then the new information is constructed with respect to the head of the rule. A IR is defined as:

$$H \longleftarrow B_1, ..., B_n, \phi_1, ..., \phi_m$$

H and $B_1, ..., B_n$ are templates; $\phi_1, ..., \phi_m$ are expressions. Expressions are formulated over the variables in the templates and constrain the possible substitution for the variables. The example template described above would serve as the head of such an integration rule, because the information from `CORINE` has to be translated into this format. As the corresponding information in `CORINE` is stored in two different tables we have two templates in the rule body, namely:

```
<clc_ns2,?LA,complex, {
  CLC_NS2_ID -> <CLC_NS2_ID,?LI,string,?ID>,
  AREA ->       <AREA,?LS,int,?S>,
  ns ->         <ns,?LC,int,?C>
  }>@CORINE,
<clc_ns2_pol,?LP,complex, {
  CLC_NS2_ID -> <CLC_NS2_ID,?LI,string,?ID>,
  VERT_ID ->    <VERT_ID,?LV,int,1>,
  RWERT ->      <RWERT,?LR,real,?X>,
  HWERT ->      <HWERT,?LH,real,?Y>,
  NEXT_V_ID ->  <NEXT_V_ID,?LN,int,_>
  }>@CORINE.
```

Please note, that each context in the templates replaced by variables. The variables are bind to the context of the `CORINE` template. The result of the rule is an item with the structure of the `ATKIS` table but in the context of the `CORINE` database. The next task during the integration is to transform the items in the `ATKIS` context.

3 Context Representation with Conceptual Models

The OIL language [11] has been developed in the context of the On-To-Knowledge Project (http://www.ontoknowledge.org) as a proposal for a language for specifying and exchanging ontologies. OIL tries to provide a core set of features that have been widely accepted to be useful. OIL combines frame-based modeling primitives, reasoning facilities from Description Logics and a tight interaction with meta-data standards on the web such as RDF and XML. We used OIL to build a semantic context model of our example data by identifying a set of common properties that can be used to define a land use class (cf. [5, 14]).

We assume that each information source has its own ontology but import their primitives from several global ontologies []. In our example the central global ontology describes areas. The ranges of the different slots are defined in several ontology modules imported by the area ontology. These ontologies define generic taxonomy of ground types, artificial structures, plants and types of use.

class-def area
  slot-constraint ground
    value-type ground-type
  slot-constraint coverage
    value-type stucture
  slot-constraint cultivation
    value-type plant
  slot-constraint vegetation
    value-type plant
  slot-constraint use
    value-type use-type

Based on the slots provided and the taxonomy of potential fillers, we built an ontology of use types for each of the catalog systems, e.g. the `CORINE` and the `ATKIS` catalog. These ontologies import the generic area ontology and define the specific land types as defined subclasses that are inherited form the area concept.

The different concepts in the two ontologies are described by characteristic slot constraint corresponding to the informal description of the land use types in the catalog definitions. The result of the modeling effort is a semantic characterization of terms that are used to indicate the type of a certain entry in a catalog system.

In our case study, which is only a small but representative subset, the overall model defining the semantics of the two information sources contained about 60 primitive concept definitions most of which were used to describe taxonomy of fillers for the 8 role definitions used. The definition of the land use classes in the two catalog systems consists of approximately 50 defined concepts each using some of the 8 role definitions with corresponding role-constraints. For example:

class-def broad-leaved-forest
  subclass-of surface
  slot-constraint coverage
    value-type no-stuctures
  slot-constraint ground
    value-type land
  slot-constraint vegetation
    value-type OR trees, shrubs
    value broad-leaved-trees

The above example is taken form one of the entries in our CORINE data-set used in the case study. This entry is classified as 'broad-leaved-forst' which is a subclass of the CORINE concept 'forest' mainly consisting of broad-leaved trees. We get a description of this concept by adopting the definitions of the super-classes 'forests' and 'forests-and-semi-natural-areas' and specializing the 'has-value' constraint on the 'vegetation' slot from 'trees' to 'broad-leaved-trees'.

4 Context Transformation

A conceptual model of the context of each information source builds a basis for an integration on the semantic level. We call this process context transformation, because we take the information about the context of the source (in our case `CORINE`) and re-interpret this information in the terms of a target (`ATKIS`) providing a new context description for that entity within the new information source. We compare two different approaches for context transformation namely rule-based context transformation already implemented in the MECOTA system and context transformation based on classification and show how these two approaches can be used to integrate the example data.

4.1 Context Transformation with Rules

Syntactically Context Transformation Rules (CTR's) are built up similar to integration rules. They also use template and expressions. But their semantics differ in two points. First a context transformation rule has *exchange* informations. Therefore CTR's replace one template by another. Second, CTR's can also be applied to template which is nested in the structure of a top–level template i.e. to an template in an attribute. A CTR is represented as follows:

$$H \longleftarrow \overline{B} \& B_1, ..., B_n, \phi_1, ..., \phi_m$$

In the CTR one template $\overline{B}$ of the body templates is needed to be signed. The signed template replaces the head template H of the rule. The other templates in the body $B_1, ..., B_n$ and the expressions $\phi_1, ..., \phi_m$ are required to support or to restrict the replacement. The inference principle of CTR's is described in detail in [20, 21].

We illustrate the power of CTR's in our example. The surfaces in `ATKIS` and `CORINE` are stored with different measures of size, namely square-meters and hectares. Therefore the surface value of `CORINE` can not be copied but has to be converted dividing the number of square-meters by the factor 10000 . The conversion is done during the context transformation. The appropriate CTR looks like:

```
<?N,surface-ATKIS,?T, ?V_M2>@?S :-
  <?N,surface-CORINE,?T, ?V_HA>@?S
  &
  ?V_M2 := ?V_HA / 10000.
```

In difference to IR's, now the contexts are specified but the name, type, value, and source are replaced by variables. Such a CTR is applicable to all templates with the specified context including those templates which are hidden in the structure, e.g. to the template `AREA` stored in the attribute `AREA` of template `clc_ns2`.

Context transformation rules do not cover all kinds of semantic heterogeneity. While we were able to integrate the diverging measures used to determine the size of area, the integration of the different kinds of type information is more difficult and demands for more sophisticated mechanisms for context transformation.

4.2 Context Transformation by Classification

Classes identify common properties of their members by defining necessary conditions for a membership. A classification problem is characterized by the determination of membership relations between an object under consideration and a set of predefined classes. The identification process starts with data about the object that has to be classified. This data is provided by so-called observation. In the course of the classification the observed data is matched against the necessary conditions provided by the class definitions leading to one or more classes. The match between observations and membership conditions is performed using knowledge that associates properties of objects with their class. This view on classification can be formalized in the following way [18]:

- Let C be a set of solution classes (in our case concept predicates $\{c_1, \dots, c_m\}$)
- Let O_2 be a set of implicit Observations (in our case the necessary conditions for class membership)
- Let R be a set of classification rules (in our case sufficient conditions for class membership)

Then in principle a classification task is to find a solution class $c_i \in C$ in such a way, that

$$O \wedge R \Rightarrow c_i(X) \tag{1}$$

In terms of the definitions given above, semantic translation is equivalent to a re-classification of entities already classified in one semantic structure $S = \{c_1^S, \cdots, c_n^S\}$ using another semantic structure $T = \{c_1^T, \cdots, c_m^T\}$. The process of re-classification can be based upon the semantic characterizations given by both structures [19]. While the definitions in the source structure S can be used to infer properties of an entity, the semantic characterizations of concepts in the target structure T defines the goal that has to be proven to classify an entity into an existing concept in T.

In order to perform this kind of classification, we translated the OIL definition given above into the description logic [16, 2, 3] underlying it, which only requires syntactic replacements. Once we have made the translation we can use the description logic reasoner FaCT [10] to compute subsumption relations between concepts from the two catalogue systems and the translated specification (cf. [4]) what is shown below:

```
(defprimconcept broad-leaved-forest (and surface
        (all coverage no-structures)
        (all ground land)
        (all vegetation (or trees shrubs))
        (some vegetation broad-leaved-trees)))
```

The derived subsumers from the source hierarchy are FORESTS-AND-SEMI-NATURAL-AREAS and FORESTS (direct subsumer). In the case of 'broad-leaved-forest' we also get the correct result for the target hierarchy. The subsumers from the target hierarchy are: VEGETATION-AREA and FOREST-AREA (direct subsumer). Inspecting the target hierarchy, we can see that this is exactly the position we would expect. So, we can say that at least for this case the context transformation problem could be solved in a straight-forward way using OIL and the FaCT reasoner. We would take the direct subsumer in the target hierarchy (forest area) to indicate the type of the translated entity in the new context of the ATKIS catalogue.

4.3 Problems with classification–based context transformation

In our case study most of the objects are re-classified correctly. But some other examples in the case study illustrate the limit of a description logic reasoner (e.g. FaCT) used for the re-classification task. One striking problem occurs when the reasoner either finds none or more than one subsumer in the target. In this case it is not clear what to do: for example, in our case study the concept DISCONTINUOUS-URBAN-FABRIC has two direct subsumers, namely RESIDENTIAL-AREA and VEGETATION-AREA.

```
(defprimconcept discontinuous-urban-fabric (and area
         (all use living)
         (all coverage (or road-network buildings))
         (all vegetation natural-plant)
         (all cultivation no-plants)
         (all ground land)))
```

For the classification one of the two direct subsumers has to be selected. In practice, DISCONTINUOUS-URBAN-FABRIC should be mapped to RESIDENTIAL-AREA, because it is mainly an area with a low density of buildings and road-networks (i.e. a RESIDENTIAL-AREA) but *sometimes* it also contains a few natural plants. Because of the presence of vegetation, the description logic reasoner derives a subsumption by VEGETATION-AREA. This example demonstrates the weakness of the pure logic–based description formalism which cannot express uncertainty.

This case leads to two important observations. First, heuristic or statistical classification mechanisms are needed providing some information about which of the results is to prefer or which subclass relation is the most likely one even if we fail to prove it. These approaches often rely on a much less expressive model while showing greater flexibility in their reasoning service.

Second, this case also displays the weakness of the rule-based context transformation, i.e. the demand for a more sophisticated context transformation than a rule–based mechanism. In principle, the result of a re-classification can be easily transformed into a simple context transformation rule. The re–classification is used in the rule acquisition task in order to acquire th appropriate rules for the context transformation. The rule contains the compilation of the classification result. But please note that this rule set of simple context transformations would be very large. In practice, such a large set is not easy to maintain. Second, in those cases where one subsumer cannot be determined a heuristic mechanism (i.e. uncertainty) has also to be combined with the rule mechanism.

In both points discussed above, more sophisticated methods for the context transformation are needed which base on heuristic mechanism. A hybrid approach integrating various methods is needed. An important research question is if they could be generalized to a common model of context transformation for semantic interoperability covering different classification approaches (rule inference, subsumption reasoning, heuristic classification, statistical classification). This topic seems to be one of the most challenging questions to be addressed by research on semantic information integration.

But in our opinion, which is substantiated by our case study, — both approaches for context transformation presented in this paper, i.e. the context transformation rules for the functional context transformation and the re-classification, are suitable for reconciling most of the semantic integration problems.

The next example demonstrates the limit of the re–classification approach itself. The underlying basic problem is that a concept needs two different concept desriptions depending on in which subsumption taxonomy the concept has to be classified. In the case study a MINERAL-EXTRACTION-SITE belongs to the artificial surfaces. These super-classes provide a first definition of the concept:

```
(defprimconcept mineral-extraction-site
    (and area
        (all coverage structure)
        (all cultivation no-plants)
        (all ground land)
        (all use extraction)))
```

MINERAL-EXTRACTION-SITE is derived as a direct subsumer of AREA in the source taxonomy `CORINE` which is explicitly stated in the definition. Unfortunately, this description completely failed to produce the appropriate result in the target hierarchy of the `ATKIS` catalogue. In order to place this concept correctly in the target hierarchy we have to redefine the USE relation to a exists statement (i.e. change `(all use extraction)` to `(some use extraction)`). The concept is now positioned correctly in the target hierarchy. However, no subsumers could be found in the source hierarchy!

This result is really negative because it provides strong evidence that specification of the information entity was strongly biased by the solution we wanted to produce. It is not clear at all whether the other results have to be judged under the same assumption.

5 Discussion

In this paper we described how a real-life information integration problem could be solved using a combination of structural and semantic integration methods. Thus, the structural integration using the rule-based mediator MECOTA served as a basis for the application of two semantic integration methods, namely rule-based and classification based context transformation. We showed that the rule-based approach is suitable for rather simple transformations like the conversion of diverging measures. We argued that more complex context transformations can be performed using a classification approach and gave evidence for this claim from a recent case study.

Remaining Problems: Our case study showed that the use of the OIL language and the FaCT reasoner enables us to perform a semantics-preserving translation of attribute values from one data source to another. However, we encountered many open problems, most of which were the result of missing tool support for the translation process, but there are also some principle problems. The most striking problem occurs when the reasoner either finds none or more than one subsumer. In this case it is not clear how to proceed. Such examples clearly show the demand for heuristic or statistical classification mechanisms combined with the reasoning services of a description reasoner.

Hybrid Transformation Approaches: Another question left open in this paper is, how the rule-based, the classification-based transformation approach, and further heuristic methods can be integrated into a common framework which can be operationalized by a mediator system. It seems quite clear that the semantic part of the integration process has to be based upon the structural integration in order to have a shared vocabulary for the transformation task. However, it is not clear at all how rule-based, classification-based, and heuristic-based transformations relate to each other. An important research question is if they could be generalized to a common model of context transformation for semantic interoperability. In the light of the problems mentioned above, we can even broaden the claim for a hybrid transformation approach by claiming that not only different general paradigms like rule-based transformation for functional relations and classification-based transformation for non-functional dependencies, but also different classification approaches (subsumption reasoning , heuristic classification, statistical classification) should be covered by a general framework. This topic seems to be one of the most challenging questions to be addressed by research on semantic interoperability.

References

[1] AdV. Amtliches Topographisch Kartographisches Informationssystem ATKIS. Technical report, Landesvermessungsamt NRW, Bonn, 1998.

[2] F. Baader. Logic-based knowledge representation. In M.J. Wooldridge and M. Veloso, editors, *Artificial Intelligence Today – Recent Trends and Developments*, volume 1600 of *Lecture Notes in Artificial Intelligence*, pages 13–41. Springer Verlag, 1999.

[3] F. Baader, H.-J Bürckert, B. Hollunder, A. Laux, and W. Nutt. Terminological logics. *Künstliche Intelligence (KI)*, 3, 1992. (in german).

[4] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Daniele Nardi, and Riccardo Rosati. Description logic framework for information integration. In *Proc. of the 6th Int. Conf. on the Principles of Knowledge Representation and Reasoning (KR'98)*, pages 2–13, 1998.

[5] Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Daniele Nardi, and Riccardo Rosati. Knowledge representation approach to information integration. In *Proc. of AAAI Workshop on AI and Information Integration*, pages 58–65. AAAI Press/The MIT Press, 1998.

[6] EEA. Corine land cover. Technical guide. Technical report, European Environmental Agency. ETC/LC, European Topic Centre on Land Cover, 1997-1999.

[7] P. Frankhauser, M. Kracker, and E. Neuhold. Semantic vs. structural resemblance of classes. *SIGMOD Record*, 20(4), December 1991.

[8] Hector Garcia-Molina, Joachim Hammer, Kelly Ireland, Yannis Papakonstantinou, Jeffrey Ullman, and Jennifer Widom. Integrating and accessing heterogeneous information sources in tsimmis. In *Proceedings of the AAAI Symposium on Information Gathering*, pages 61–64, Stanford, California, March 1995.

[9] Hector Garcia-Molina, Yannis Papakonstantinou, Dallan Quass, Anand Rajaraman, Yehoshua Sagiv, Jeffrey Ullman, and Jennifer Widom. The tsimmis approach to mediation: Data models and languages. In *Next Generation Information Technologies and Systems (NGITS-95)*, Naharia, Israel, November 1995. Extended Abstract.

[10] I. Horrocks. The FaCT system. In H. de Swart, 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, May 1998.

[11] I. Horrocks, D. Fensel, C. Goble, F. van Harmelen, J. Broekstra, and M. Klein. The ontology interchange language oil: The grease between ontologies, to appear.

[12] V. Kashyap and A. Sheth. Schematic and semantic similarities between database objects: A context-based approach. Number TR-CS-95-001, LSDIS Lab, Univ. of GA, 1995.

[13] Vipul Kashyap and Amit Sheth. *Cooperative Information Systems: Current Trends and Directions*, chapter Semantic Heterogeneity in Global Information Systems: The role of Metadata, Context and Ontologies. 1997, Academic Press.

[14] Zoubida Kedad and Elisabeth Métais. Dealing with semantic heterogeneity during data integration. In Jacky Akoka, Mokrane Bouzeghoub, Isabelle Comyn-Wattiau, and Elisabeth Métais, editors, *Conceptual Modeling - ER '99, 18th International Conference on Conceptual Modeling*, volume 1728 of *Lecture Notes in Computer Science*, pages 325–339. Springer, November, 15-18 1999.

[15] W. Kim and J. Seo. Classifying schematic and data heterogeinity in multidatabase systems. *IEEE Computer*, 24(12):12–18, 1991.

[16] M. Schmidt-Schau"s and G. Smolka. Attributive concept descriptions with complements. 48(1):1–26, 1991.

[17] Edward Sciore, Michael Siegel, and Arnie Rosenthal. Using semantic values to facilitate interoperability among heterogeneous information systems. *ACM Transactions on Database Systems*, pages 254–290, June 1994.

[18] Mark Stefik. *Introduction to Knowledge Systems.* San Francisco, California, 1995.

[19] Heiner Stuckenschmidt and Ubbo Visser. Semantic translation based on approximate re-classification. In *Workshop on Semantic Approximation, Granularity and Vagueness*, Breckenridge, Colorado, 2000.

[20] Holger Wache. A rule-based mediator for the integration of heterogeneous sources (extended version). TZI–Berichte, University of Bremen, Bremen, 1999.

[21] Holger Wache. Towards rule–based context transformation in mediators. *Proceedings of the International Workshop on Engineering Federated Information Systems (EFIS 99)*, 1999.

[22] Gio Wiederhold. Mediators in the architecture of future information systems. *IEEE Computer*, 25(3):38–49, March 1992.